

Choosing a Set of Coherence Relations for Text Generation: A Data-Driven Approach

Alistair Knott¹ and Robert Dale²

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

² Robert Dale, Microsoft Institute, 65 Epping Road, North Ryde, NSW 2113, Australia. Email: rdale@microsoft.com

Abstract. An active research programme in Natural Language Generation has grown up around the notion of ‘coherence relations’. Relations are being used in a variety of roles in generation systems, in planning the structure of messages and in guiding their linguistic realisation. However, a confusing diversity exists between systems with respect to the particular relations they use.

This paper outlines a systematic methodology for justifying a set of relations, on the basis of the **cue phrases** that can be used to signal them explicitly in text. A large taxonomy of cue phrases has been gathered, from which a conception of relations as feature-based constructs is motivated. By considering small extracts from the taxonomy, individual features of relations can then be investigated in isolation. The paper concludes by noting some advantages of a feature-theoretic treatment of relations for NLG, as well as some problems with a reliance on cue phrases in a theory of text coherence.

1 Introduction

In the field of natural language generation, the notion of **coherence relations** has recently found widespread application. Relations feature in a number of computational theories of discourse coherence (Hobbs [9], Grosz and Sidner [8], Mann and Thompson [20]), and have been adapted for use in NLG with considerable success.

But the popularity of coherence relations in natural language generation has given rise to its own problems: researchers have altered the original relations to suit their own needs, and as a consequence there is an increasing diversity amongst the relations which are being used. This proliferation raises the question of how to decide on a particular set of relations for text generation—a question which has been addressed by several researchers already, and to which a new answer is suggested in this paper.

We will begin in Section 2 by outlining Mann and Thompson’s conception of **rhetorical relations**, which has proved particularly influential in NLG. We then review the different ways in which rhetorical relations have been adapted for the task of text generation. In Section 3 we will outline some of the problems associated with the diversity of sets of relations currently in use; and in Section 4

we discuss some of the sources of justification for relations which have been drawn on (either explicitly or implicitly) by researchers in the field. Two of these are of particular interest: one is the suggestion that relations should be thought of as ‘psychologically real’; the other is the intuition that relations should be linked to surface linguistic phenomena. In Section 5, we argue that these two sources can be combined, and that linguistic phenomena (in particular **cue phrases** such as *because* and *nevertheless*) can be used as evidence in determining the set of ‘psychologically real’ relations. On the basis of this argument, a methodology is proposed in Section 6 for justifying a set of relations by collecting cue phrases, arranging them into a taxonomy of synonyms and hyponyms, and using this taxonomy to motivate relation definitions. Section 7 outlines some advantages of the new conception of relations, and Section 8 discusses some of its potential problems.

2 Coherence Relations in NLG Systems

2.1 Rhetorical Structure Theory

The central contention of Mann and Thompson’s RST³ is that a text is coherent mainly in virtue of the **rhetorical relations** which link together its various sub-parts. The claim is that any coherent text can be broken down into a **text spans** of different sizes, and that these spans can be linked to one another by a hierarchical arrangement of rhetorical relations. A sample RST analysis is given in Figure 1.

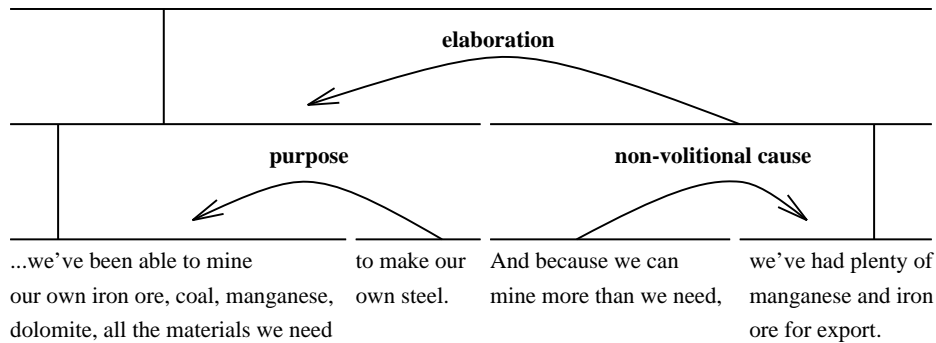


Fig. 1. A Sample Rhetorical Structure Tree (adapted from Mann and Thompson [20])

Mann and Thompson claim that a set of 23 relations suffices to analyse ‘the vast majority’ of English prose texts. This set includes relations such as ELABORATION, CONTRAST and JUSTIFY.

An important feature of RST is that relation definitions make no reference to surface structures in text; the applicability of a relation in a given situation is

³ By ‘RST’ we mean that version of the theory presented in Mann and Thompson [20].

determined entirely by the judgements of a discourse analyst. Textual devices for ‘signalling’ relations (for instance connective words like *because* or *but*) are not always found; and even when found, do not stand in a one-to-one mapping with relations—so they cannot be relied upon when carrying out an RST analysis.

2.2 ‘Operationalising’ Rhetorical Relations

RST was conceived as a static, *descriptive* theory of textual coherence. But rhetorical relations have since been adapted to *computational* roles in a number of generation systems. Hovy’s original text structurer [11] [10] makes use of relations as STRIPS-style **planning operators**. Moore and Paris’ generator ([24] [23] [25] [26]) also uses relations in planning operators, but also features operators associated with higher level intentions.

Attention has also been focused on the linguistic realisation of rhetorical relations. Initial systems such as Hovy’s associated relations with ‘typical’ connective phrases (e.g. *then* and *next* for the SEQUENCE relation) or with grammatical devices (e.g. an embedded relative clause to mark ELABORATIONS). Several such techniques were discussed by Scott and Souza [31], who presented a number of heuristics for guiding the textual realisation process. More recently Rösner and Stede [27] [28] and Knott [15] make the realisation of a relation sensitive to its hierarchical position in the rhetorical structure tree—Rösner and Stede also allow for ‘elegant variation’ in the way relations are marked. Some researchers (notably Vander Linden *et al* [34] [33] and Rösner and Stede) have also looked at ways to extend RST analysis within single clauses.

Finally, some NLG systems envisage multiple relations applying simultaneously between two spans of text, representing different aspects of its meaning. Moore and Pollack [22] argue that generation systems require a representation of both INTENTIONAL and INFORMATIONAL relations; Hovy *et al* [13] and Maier (this volume) use three levels of relations in their system, representing IDEATIONAL, INTERPERSONAL and TEXTUAL information. Elhadad and McKeown’s [4] algorithm for generating connectives operates with relations on four different dimensions; ARGUMENTATIVE ORIENTATION, FUNCTIONAL STATUS, POLYPHONIC FEATURES and THEMATIZATION.

3 The Proliferation of Relations

Despite the attractions of the general concept of rhetorical relations in text generation, there is a surprising amount of variability between researchers in the sets of relations they use. Virtually no two systems use exactly the same set of relations; and few systems use exactly those relations proposed in RST.

Departures from RST’s original set are of many different types. RST relations have been subdivided—for instance, Rösner and Stede’s STEP-SEQUENCE is a specialisation of the RST relation SEQUENCE, as is Kosseim and Lapalme’s [18] ACTION-SEQUENCE. RST relations have also been amalgamated to form new relations; for instance, Scott and Souza combine the relations VOLITIONAL-CAUSE,

NON-VOLITIONAL-CAUSE, VOLITIONAL-RESULT, NON-VOLITIONAL-RESULT and EVIDENCE, for the purposes of textual realisation. Alternative cuts can also be found through the space of relations—for instance, Hovy *et al*'s [13] MEANS seems to incorporate aspects of the RST relations ELABORATION and CIRCUMSTANCE; Delin *et al*'s [3] GENERATION and ENABLEMENT likewise cut across RST's existing relations. Finally, some new relations have appeared which have no obvious connection with RST relations at all—for instance, ALTERNATIVE (used by Scott and Souza, and Rösner and Stede). In short, a large number of very different sets of relations are currently in use in NLG.

As the sets of relations employed by researchers become more diverse, problems begin to appear concerning the explanatory power of the relations being used. Relations were originally elements in a theory of discourse coherence, which stated that 'virtually any text' can be analysed by representing its rhetorical relations. But if new relations can be created whenever they are needed, this claim at best loses much of its strength, and at worst becomes unfalsifiable.

To avoid this problem, a way is needed of determining *in advance* what is to count as a relation, and then of ensuring that all the relations in a given system conform to this specification.

In other words, a set of relations needs to be justified. This justification cannot simply be an observation that the set of relations in question suffices to analyse a given class of texts, no matter how extensive this is. The *claim* about relations is that they can be used to analyse texts—an independent way of characterising them is needed if they are to do any empirical work. Any number of sets of relations could be chosen to describe texts; for instance the pair 'CAUSAL' and 'NON-CAUSAL' would work. Again, we could describe texts using relations like 'MALE-VOLITIONAL-CAUSE' and 'FEMALE-VOLITIONAL-CAUSE'—as a 'purely descriptive' mechanism, this cannot be faulted. Of course, we want to do more than just describe a text with a set of relations. We want the description to contribute towards an explanation of why it is or is not coherent.

4 How Can a Set of Relations be Justified?

RST claims to be a 'purely descriptive' theory of text structure. But although it does not seek to justify its set of relations, there is intuitively something plausible about them (and about many of the 'new' relations also), which explains why a relation like 'PURPOSE' is included, and not a relation like 'FEMALE-VOLITIONAL-CAUSE'.

Two possible motivations for relations seem to be implicit in Mann and Thompson's theory. One is that relations are somehow part of our **psychological makeup**—we use them when we construct and interpret text. This is never stated outright in RST; but it seems intuitive that we are thinking about things like 'contrasts', 'sequences' and 'justifications' when we process text. It is equally intuitive that, when we are analysing the discourse relations in a text, we are not normally concerned about whether the agent involved in a causal relation is male or female.

A second plausible source of motivation for RST relations is the existence of certain linguistic phenomena which appear to correlate with them quite closely. Mann and Thompson’s relations are clearly not defined in terms of linguistic phenomena. But this is not to say they do not *mirror* a class of linguistic phenomena. The very fact that algorithms can be devised for the ‘linguistic realisation’ of relations is an indication to the contrary.

Both of these sources of motivation can be used to give an independent definition of relations. The following two sections outline some of the more explicit justifications of relations that have been given on these grounds.

4.1 Relations Conceived as ‘Psychologically Real’

The notion that relations model psychological constructs finds its clearest expression in the work of Sanders *et al* [29] [30]. They propose a parameterisation of the space of relations using four ‘cognitively basic’ primitives—categories like **causality** and **polarity**—and seek out empirical evidence for their proposal in psychological experiments.

The assumption of psychological reality gives relations a clear theoretical role. We are interested in what makes a text coherent: it is likely that this explanation will make reference to the way texts are actually constructed and interpreted. Texts are, after all, created by and designed for humans, so we should expect an explanation of coherence to refer to human processing mechanisms.

But the empirical question of deciding *which* relations people use is a hard one to answer satisfactorily. Sanders *et al*’s experiments give better evidence for some parameters than for others; and their grounds for deciding which parameters to test for in the first place are also debatable (see Knott and Dale [17]). In addition, Sanders’ four parameters only determine *categories* of relation; they suggest further distinctions within these categories which are not motivated by experiment. It may be that psychological experiments on readers and writers are too blunt a tool with which to investigate the fine detail of a set of relations—the sort of detail which may be needed for the purposes of text generation. For one thing, subjects’ introspections about the relations they are using cannot be guaranteed: the processes are not open to conscious inspection. And yet it is hard to extrapolate from subjects’ observable behaviour to the relations they are using: the experimental windows we have (for instance reading time, or accuracy of recall) tend to be extremely indirect, and are likely to be sensitive to other factors for which it is difficult to control.

4.2 Relations Associated with Linguistic Phenomena

Many researchers (eg. Hobbs [9], Rösner and Stede, Sanders *et al*) associate relations with ‘typical’ **cue phrases**—phrases like *because* and *however*. Hovy [12] appears to take a further step, using cue phrases as *evidence* for a taxonomy of relations: the fact that the word *then* can mark a temporal and a spatial sequence, but that *after* and *beside* are specific to temporal and spatial sequences

respectively, is an argument for classifying these two relations as subordinate to a general SEQUENCE relation.

Cue phrases can thus provide a neat way of classifying relations. Moreover, it should be possible to justify a fairly detailed taxonomy using this source of evidence. But the project of justifying a taxonomy using cue words is far from complete: Hovy does not go into any detail, and his later classification (Hovy *et al* [13]) is not systematically motivated from cue phrases. More importantly, an account is needed as to *why* cue phrases should be considered evidence for relations. Hovy admits that his evidence is ‘nonconclusive’; it seems clear that relations thus conceived are not doing the kind of explanatory work which psychological constructs do. How can we use the association with cue phrases to *justify* our taxonomy over alternative taxonomies where such associations are not made? Why should the association with cue phrases be particularly significant? Linking relations to cue phrases is convenient, but it is still not clear if it is of theoretical relevance.

5 Cue Phrases: Linguistic Evidence for Cognitive Constructs

If we could think of relations as modelling psychological constructs, *and* use cue phrases as empirical evidence for these same constructs, then the drawbacks of both the above approaches can be minimised. The assumption of psychological reality ensures that relations have a genuine explanatory role in a theory of text coherence; and the use of cue phrases as a source of evidence avoids the need to rely on psychological experimental approaches, which are awkward for justifying a detailed taxonomy. Clearly, it would be advantageous to combine these sources of justification. Furthermore, there are good reasons—outlined below—to suppose that they *can* be combined.

5.1 A Conception of The Role of Relations in Human Text Processing

The argument that cue phrases are likely to reflect ‘the relations we use when we process text’ is based on a particular conception of what these relations are, and how they work. The essential idea is that relations, expressing constraints on the formation of coherent texts, are used to *reduce the search spaces* involved in human reading and writing tasks. For example, a writer, in the middle of producing a text, knows that the next span to be produced must be linked to the previous one by one of the coherence relations; so coherence relations can be used as search keys for relevant content. Likewise, a reader in the middle of a text could use the set of relations to create hypotheses about the next span to be processed, and so facilitate its processing.

This conception of relations is plausible from the point of view of computational theory: writing and reading are very complex tasks, hence any strategies

for reducing this complexity would be highly desirable. Empirical evidence for relational constructs can also be found. For instance, Bereiter and Scardamalia [1] (p.62) have shown that prompting child writers with phrases like *even though* or *for example* seems 'to stimulate the children to search for new nodes in memory' when they have 'dried up'. It could be that in more mature writers, relational constructs take the place of such external cues, so that they can initiate the search for content themselves. As regards reading, there are many experiments which show that the existence of relations in a text facilitates its processing. For instance, Meyer and Freedle [21] have shown that a text where the elements of content are linked by causal relations is recalled better than one where the same elements are presented but without such relations to link them.

5.2 The 'Communication of Relations'

In the previous section, a conception of relations was advanced in which they mediated the tasks of both reading and writing. It was proposed that writers generate a structure of relations as an intermediate step in creating a text; and that readers build a structure of relations as an intermediate step in understanding a text. If this idea is accepted, then we can think of the *communication* of relations as an intermediary step in the overall process of communication via a text: the reader can be thought of as working out the relations which the writer has used. In other words, relations *really matter*: it is important for the writer to make the relations in a text clear to the reader.

It is for this reason that we can expect there to be linguistic ways of *signalling* those relations which are involved in the processes of writing and reading.⁴ Of course, a writer will not always have to signal a relation explicitly: often, it will be inferrable from context, or from the reader's world knowledge. However, it is unlikely that any relation exists which would *always* be thus inferrable. The 'inferrability' of a relation is not something intrinsic to the relation itself, but rather something which depends on the whole communicative situation, including context and the reader's background knowledge. And these are things which will clearly vary from text to text. Thus it still seems likely that for every 'psychologically real' relation, there exists a corresponding linguistic device for making that relation explicit.

5.3 Why Look at Cue Phrases?

The final question to be addressed is, what constitutes a 'linguistic device for making relations explicit'? The problem is that we do not as yet have *definitions* of the different relations—these are what we are working towards—so it is not immediately obvious how to identify *markers* of relations. In the present work,

⁴ Note that if we did not think of relations as psychological constructs, this argument would not go through. If relations are conceived just as descriptive devices, we have no reason to suppose they will be marked in language: why should they be if they are just being used by analysts to describe texts?

we have chosen to concentrate on ‘cue phrases’—that is, on sentence and clause connectives. They seem the obvious place to begin, because they so clearly have an import beyond the clause or sentence in which they are found. What is more, many cue phrases have their origins in more complex formulas and phrases—for instance *because* originally contained an anaphoric expression and an explicit reference to a cause; *however* has evolved from expressions like ‘however that is’ or ‘however that might be’. The phrases have clearly been shortened and standardised because the concepts they represent are frequently made use of, just as we would expect relations to be.

Despite such considerations, it is still likely that relations can be textually marked by other means than cue phrases, and these should be considered in future work.

6 A Data-Driven Methodology for Justifying a Set of Relations

We have argued that cue phrases can be used as evidence for relations. The following sections outline the methodology which we have developed and used to this end, and describe the results which were achieved.

6.1 A Test for Cue Phrases

A first task was to give an initial definition of the term ‘cue phrase’. Many researchers who use this term (eg Grosz and Sidner [8]) simply provide a list of examples to illustrate its application, or define it within the terms of the theory they are proposing—we wanted a theory-neutral way to pick out *any* cue phrase, and thus have formulated a linguistic test for doing so. The test is designed to capture the idea that cue phrases are phrases which link clauses together, and cannot function in the context of a single clause. It works as follows:

1. Choose a phrase in a piece of naturally-occurring text: the **candidate phrase**.
2. Extract the clause in which the phrase appears, and consider it in isolation.
3. Substitute any anaphoric or cataphoric terms in the clause with their antecedents, and include any elided items.
4. If the candidate phrase is a cue phrase, the resulting text should appear *incomplete*; i.e. it should require additional linguistic context in order to be interpreted. For instance, Text 1 is incomplete, while Text 2 is complete.⁵
 - (1) *Nevertheless*, Sam stayed at the beach.
 - (2) *Yesterday*, Sam stayed at the beach.
5. If the candidate phrase is then removed, the remaining clause should *no longer* be incomplete. In other words, it should form a coherent ‘mini-discourse’ on its own. If the word *nevertheless* is removed from Text 1, for example, the resulting text can stand on its own:
 - (3) Sam stayed at the beach.

⁵ For more detail about this test, see Knott and Dale [17].

6.2 Gathering a Corpus of Cue Phrases

A substantial corpus of cue phrases was gathered using this test. To limit the scope of our investigation, the texts we used were all from academic articles and books. In all, 226 pages of text were analysed, from 12 different authors: this yielded a corpus of around 200 different phrases.

Cue phrases were found to fall into a variety of syntactic groups. **Coordinators** such as *and* and *but* appear between the clauses they link; **subordinators** such as *because* and *although* introduce subordinate clauses in complex sentences; **conjunct adverbs** such as *however* and *above all* modify single clauses, and can appear at different points within them; **phrases taking sentential complements** include phrases like *it follows that* and *it may seem that*.

6.3 A Test for Classifying Cue Phrases

The next stage was to organise cue phrases according to their function in signalling relations. Again, a theory-neutral method was sought; classification was based on a second linguistic test, this time to investigate whether one phrase is **substitutable** for another in a given context. The task of substituting cue phrases is one which occurs frequently during the course of normal writing—the idea is to capture the judgements of ordinary readers and writers, rather than appealing to explicitly ‘theoretical’ ideas about the meaning of different cue phrases. The test works as follows:

1. Take any cue phrase from the corpus and place it in a text where it naturally occurs.
2. Imagine you are a writer who has just written this text, but needs to choose a *different* cue phrase to replace it (perhaps for reasons of elegant variation). Which other cue phrases would you consider appropriate?

An example of the substitutability test is shown in Text 4.

$$(4) \quad \text{Jim wanted to make a big } \left\{ \begin{array}{l} \textit{so} \\ \checkmark \textit{therefore} \\ * \textit{although} \\ * \textit{on the contrary} \end{array} \right\} \text{ impression that night; } \left. \begin{array}{l} \text{he took a big bunch of} \\ \text{roses with him.} \end{array} \right\}$$

In this notation, *so* is the original cue phrase; *therefore* is shown as substitutable for *so* in this context, and *although* and *on the contrary* are not. Note that the issue at stake is not whether the resulting text is grammatically well-formed—each of the cue phrases in this example is acceptable from this point of view—but rather whether a writer would substitute a given phrase for the original phrase *so*. Note also that some differences between cue phrases are ignored; for instance *therefore* belongs to a more formal register of discourse than *so*. Again, see Knott and Dale [17] for more details.

Generalising over all contexts, there are four possible substitutability relationships between two phrases *X* and *Y*.

- *X* is **synonymous with** *Y* if in any context where one can be used, the other can also be used.

- X and Y are **exclusive** if they can never be substituted for one another in any context.
- X is a **hypernym of Y** —and Y is a **hyponym of X** —if whenever Y can be used, so can X ; but there are some contexts where X can be used and Y cannot.
- X and Y are **contingently substitutable** if there are some contexts where they can be substituted, other contexts where X can be used and not Y , and still other contexts where Y can be used and not X .

These relationships can be represented in diagrammatic form, as illustrated in Figure 2.




Fig. 2. Diagrammatic Representation of Substitutability Relationships

6.4 A Taxonomy of Cue Phrases

To investigate the entire range of cue phrases, it would be very useful to have a complete representation of the relationships between them. To this end, a taxonomy of cue phrases has been created, using the graphical notation illustrated in Figure 2. The taxonomy presently incorporates around 150 cue phrases, and represents the substitutability relationship between every phrase and every other phrase.

The taxonomy makes extensive use of inheritance. A cue phrase inherits all the substitutability relationships of its hypernyms; thus in Figure 3 D is a hyponym of B , so by inheritance, D is a hyponym of A and exclusive with C . Note that the an inherited contingent substitutability relationship can be overridden, as in Figure 3 (ii) where R and S should be interpreted as exclusive.

It would have been convenient if the cue phrases in the taxonomy could have been arranged into a number of clearly demarcated categories, such that the phrases from one category were exclusive with those from each other category. The top of the taxonomy would then have looked something like Figure 4, and the cue phrases in each category could have been considered in isolation. But while this modular approach would have made things much easier, no such high-level grouping was found; for every candidate grouping, there were several




Fig. 3. Inheritance in Substitutability Diagrams




Fig. 4. The Ideal Organisation of the Top Level of the Taxonomy

phrases which appeared in more than one group. Instead of a rigid high-level classification, therefore, the taxonomy was organised into a number of overlapping categories, as illustrated in Figure 5. Each category has a subset of **exclusive**

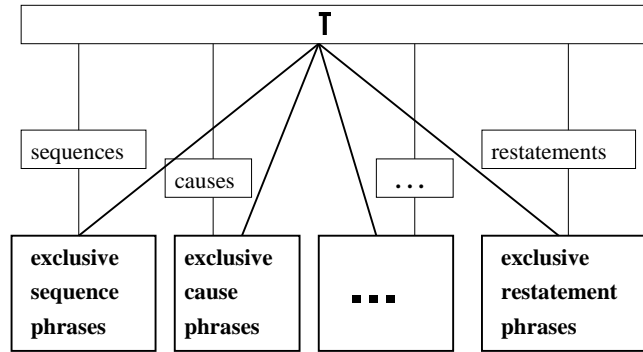


Fig. 5. The Actual Top Level of the Taxonomy

phrases (shown in bold type), which cannot be substituted for the exclusive phrases in any other category. The remainder of its phrases are **multicategory phrases**, which appear in more than one category.

The taxonomy is given in a number of separate substitutability diagrams, one for each category. So as to represent the relationship between every pair of phrases in the taxonomy, each diagram contains *all* the multicategory phrases, as well as the exclusive phrases for a single category. There are 10 categories in all,

with names like ‘sequence’, ‘cause’, ‘restatement’, ‘similarity’, and ‘digression’. No particular theoretical significance should be attached to these categories; they are simply a way of spreading the taxonomy over several pages. Most of the interesting relationships between cue phrases are represented at a much lower level of hierarchy, in the microstructure of the taxonomy.

An extract from the taxonomy is given in the Appendix, along with some motivating examples. The complete taxonomy is given in Knott [16].

6.5 A Theoretical Interpretation of the Taxonomy

The taxonomy does not express any theoretical claims about the meaning or function of cue phrases; it is simply a repository for data about how writers substitute one cue phrase for another. However, it can readily be given a theoretical interpretation.

In fact, the taxonomy lends itself well to a conception of relations as feature-based constructs, similar in structure to those proposed in GPSG (Gazdar *et al* [5]). The more general ‘hypernymic’ cue phrases in the taxonomy clearly convey less information about a relation than the more specific phrases at its leaves; we can think of the general phrases as signalling just *some features* of a relation, leaving others undefined. Indeed, each substitutability relationship can be given an interpretation according to this model.

- If X is **synonymous** with Y , then they signal the same values of the same features.
- If X is **exclusive** with Y , then they signal *different values* of at least one feature.
- If X is a **hyponym** of Y (and Y is a **hypernym** of X), then X signals all the features that Y signals, and *some other* feature(s) in addition, for which Y is undefined.
- If X and Y are **contingently substitutable**, then X and Y signal the same values of some features, but in addition X is defined for one or more features for which Y is undefined, and Y is defined for one or more features for which X is undefined.

Why are some cue phrases undefined for certain features? An explanation can be suggested by appealing to Grice’s maxim of quantity [6]—there is no need explicitly to signal the values of features which can easily be inferred by the reader. Imagine a situation where the value of one feature is easily inferrable but the values of other features are not; in this case it makes sense to use a cue phrase which only signals the values of these latter features.

Once a feature-theoretic conception of relations has been established, individual features can be motivated systematically by considering different extracts from the taxonomy. For instance, consider the extract given in Figure 6, taken from the taxonomy given in the Appendix.

In this example, *so* is shown to be a common hypernym of *it follows that*, *thereby*, *at that* and *instantly*. All these phrases thus inherit some feature signalled by *so*; we might postulate that *so* signals a CAUSAL relation, as do all its

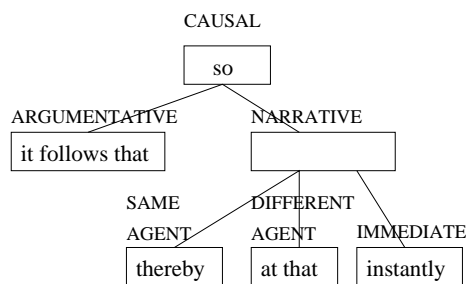


Fig. 6. Some CAUSAL Phrases

hyponyms. The diagram also shows that *it follows that* is exclusive with *thereby*, *at that* and *instantly*. We must therefore envisage a second feature for which these phrases signal alternative values; perhaps ARGUMENTATIVE for *it follows that* and NARRATIVE for *thereby*, *at that* and *instantly*. Note also that *thereby* and *at that* are exclusive phrases; we might suggest a feature with alternative values of SAME-AGENT and DIFFERENT-AGENT to account for this relationship. Finally note that both *thereby* and *at that* are contingently substitutable for *instantly*; we can thus postulate a fourth feature, which takes a value of IMMEDIATE for the phrase *instantly*, but which is undefined for *thereby* and *at that*. (Note that *instantly* is itself undefined for the feature which distinguishes between *thereby* and *at that*.)

The features proposed above should not be accorded too much importance; they are used simply to illustrate the method for motivating features from the taxonomy. Clearly, the motivation of a definitive set of features requires much more work. For one thing, a formalism for defining features must be provided, so that detailed definitions can be devised for each feature; only then can it be tested how appropriate a given feature is for describing a given cue phrase. For another thing, it is important to consider relationships throughout the taxonomy when deciding on a set of features; in the interests of parsimony, it is preferable that features do service throughout the taxonomy, rather than simply being of use in one or two cases. Finally, it must be ensured that the different features are orthogonal to one another, so that there is no redundancy in the representation they provide. A more definitive set of features adhering to these principles is currently being prepared; again, see Knott [16] for details.

7 Advantages of the New Definitions for NLG

In the previous sections a methodology for motivating a set of relation definitions has been presented which results in a conception of relations as feature-based constructs, closely tied to the cue phrases which can signal them. This section describes some of the benefits of this conception of relations for natural language generation systems.

7.1 The Textual Realisation of Relations

Firstly (and most obviously), making an explicit link between relations and cue phrases facilitates the task of signalling relations in surface text. In theories such as RST, relations are not associated in any way with linguistic devices; presumably, therefore, the task of choosing how to express a relation is a non-trivial one. In the present theory, the decision about which relation to use directly determines the set of cue phrases appropriate for signalling it. The only remaining task is to decide which features of the relation can be easily inferred by the reader, and hence left implicit.

While theorists such as Mann and Thompson are reluctant to posit a strong link between relations and cue phrases, such a link is often found in generation systems, where relations are associated with ‘prototypical’ markers. Clearly, this strategy has engineering expedience, in simplifying the algorithm for textual realisation. The present account of relations can be seen as legitimising the strategy from a theoretical point of view.

7.2 A Modular Treatment of Structural Information

As noted in Section 2.2, many generation systems envisage several different *types* of relation, tailored to capture different types of information about text structure. For instance, Hovy and Maier distinguish between IDEATIONAL, INTERPERSONAL and TEXTUAL relations; Moore and Pollack distinguish between INFORMATIONAL and INTENTIONAL relations. In each case, relations of different types can apply simultaneously between two text spans. The decomposition of relations has benefits for the architecture of a generation system, as it allows different types of information to be represented and manipulated separately.

The present methodology also envisages a decomposition of relations, this time into a set of orthogonal features. In fact, the decomposition continues beyond the two and threefold distinctions noted above; many more independent features than this are needed to represent all the relationships in the taxonomy.

Nonetheless, something very like the INFORMATIONAL/INTENTIONAL distinction emerges as one of the most productive features in the taxonomy. Consider the extract in Figure 7. All of these phrases signal a contrast of some kind. But




Fig. 7. SEMANTIC and PRAGMATIC Contrastive Phrases

while *whereas* signals a contrast between the propositional content of two spans,

then again is used when two spans are used to argue towards opposite conclusions; the contrast is here in the argumentative force of the spans, rather than their propositional content.

- (5) Kate and Sam are like chalk and cheese. Sam lives for his books; $\left\{ \begin{array}{l} \textit{whereas} \\ \checkmark \textit{ on the other hand} \\ * \textit{ then again,} \end{array} \right\}$ Kate is only interested in martial arts.
- (6) I don't know where to eat tonight. The Star of India is always good; $\left\{ \begin{array}{l} \textit{then again,} \\ \checkmark \textit{ on the other hand} \\ * \textit{ whereas} \end{array} \right\}$ we had curry just the other night.

(Note that *on the other hand* can be used in both types of situation.)

7.3 'Precondition' and 'Postcondition' Features

The decomposition of relations into features is also well suited to a conception of relations as planning operators (in which guise they frequently appear in generation systems). Planning operators are defined in terms of two components: **preconditions** (representing what must be true in order to apply an operator) and **postconditions** (representing the effects obtained by applying it).

Many of the features motivated from the taxonomy can be identified as relating solely to one or other of these components. Consider the extract in Figure 8. *So* and *however* are exclusive phrases, but both have a causal component. *So*




Fig. 8. CAUSALITY and POLARITY: 'Precondition' and 'Postcondition' Features

signals a result, as we have already seen, while *however* is used in cases where an expected result is not forthcoming:

- (7) Jim had no money for the bus, $\left\{ \begin{array}{l} \textit{so} \\ * \textit{ however} \end{array} \right\}$ he had to walk home.
- (8) Jim had plenty of money for the bus; $\left\{ \begin{array}{l} \textit{however} \\ * \textit{ so} \end{array} \right\}$ he decided to walk home.

We might envisage a feature called POLARITY, with alternative values of +VE and -VE to express this distinction.

Note that *so* and *however* both depend for their effect on the reader believing some general causal rule such as 'if X has money for the bus, X will take the bus home'. The CAUSAL feature of these phrases can thus be identified as a precondition. The difference between the phrases relates to the representation the reader ends up with; in one case the rule is obeyed, and in the other case it is broken. The POLARITY feature can thus be classed as a postcondition.

8 Problems of Descriptive Coverage: Cases where no Cue Phrase Can be Used

An obvious objection can be levelled at the methodology proposed in this paper. Its central assumption is that coherence relations can be associated with cue phrases—however, it is easy to find examples of coherent texts where no cue phrase is appropriate. Consider Text 9:

| | | | |
|-----|--|---|--|
| (9) | The Open Day Lecture was given by Mary Archer, President of the Foundation. | $\left. \begin{array}{l} * \textit{Incidentally}, \\ * \textit{In fact}, \\ * \textit{Furthermore}, \\ * \textit{Specifically}, \end{array} \right\}$ | In her talk, she discussed recent developments in cancer research. |
|-----|--|---|--|

Note that the problem in this case is not that there *is* no cue phrase in the text. This could just be due to the relation being easily inferrable. The problem is rather that no cue phrase is suitable *at all* in the text at this point, even if the writer wished to be as explicit as possible.⁶ How can the present methodology account for cases such as this?

Text 9 would be analysed by Mann and Thompson [20] with the ELABORATION relation. In fact, in the texts we have examined, the great majority of situations where no cue phrase can be used can be analysed as ELABORATIONS. According to Mann [19] (p372), ELABORATION is the relation most commonly found in text analyses; consequently the lack of a corresponding cue phrase is a serious problem.

It should be noted, however, that few theorists use relations to account for *every* aspect of text coherence. Additional mechanisms are often envisaged—in particular, that of **focus**. The metaphor of focus (Sidner [32], Grosz, Joshi and Weinstein [7], Brennan *et al* [2]) is used to represent those objects or entities in a text's domain of reference which the reader is paying particular attention to. Constraints on **focus shifts** can then be proposed, to ensure that the focus does not change too dramatically from one section of a text to the next. Many models of text processing incorporate both relations and focus; see for instance Grosz and Sidner [8] and Hovy and McCoy [14].

The metaphors of relations and focus are both very expressive, and there is a significant possibility for overlap between them. Much of the overlap can be traced to the ELABORATION relation, particularly when the elaboration is of an object or an entity; a text such as that in Example 9 might be described either as an ELABORATION or as an instance of a legal focus shift. Describing it in both ways is likely to involve a measure of redundancy.

Associating relations with cue phrases certainly does not allow relations by themselves to provide a descriptively adequate account of text. However, it does seem promising as the basis for a principled distinction between relations and focus. A model is suggested in which the coherence of a text at any point is ensured *either* by a relation *or* by focusing constraints, rather than the combined

⁶ The cue phrase *and* is perhaps marginally acceptable. But this phrase is extremely general and by no means 'identifies the relation explicitly' in this case.

action of both devices. The new model seeks to avoid some of the redundancy present in current accounts of relations and focus. Clearly, much remains to be done in working out the details of this model, and in testing it. But it is a promising direction for further research.

9 Conclusions

The central contention in this paper is that cue phrases can be used as evidence for the relational constructs which people use when they plan and interpret text. On the basis of this claim, a methodology is proposed for gathering and classifying a corpus of cue phrases—the taxonomy can then be used to motivate a set of relation definitions. This approach allows detailed relation definitions to be formulated and systematically justified, while maintaining the explanatory role of relations in a theory of text coherence. The new set of relations are feature-based constructs, with a number of advantages for natural language generation systems.

In this paper we are mainly concerned to propose a methodology for justifying relations, rather than a particular set of relations. The taxonomy we have provided might be queried in some places; likewise, different proposals might be made about the features that should be used to describe the taxonomy. But at least the method for deciding on a set of relations has been well set out, and any discussion will centre around the use of concrete linguistic examples. Indeed, one of the central advantages of the taxonomy of cue phrases is in providing a solid framework for discussions such as these.

Motivating relations from cue phrases provides at best a partial account of text coherence; there are often places in a text where no cue phrase is appropriate. We suggest that a more suitable account of coherence in these cases is provided by the focus metaphor; in fact, our conception of relations is promising as a way of reducing the overlap that currently exists between the metaphors of relations and focus. We intend to pursue this matter in further research.

References

1. C Bereiter and M Scardamalia. *The Psychology of Written Composition*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1987.
2. S E Brennan, M W Friedman, and C Pollard. A centering approach to pronouns. In *Proceedings of the 25th ACL Conference*, pages 155–162, Stanford, 1987.
3. J Delin, A Hartley, C Paris, D Scott, and K Vander Linden. Expressing procedural relationships in multilingual instructions. In *Proceedings of the 7th International Workshop on Natural Language Generation*, pages 61–70, Kennebunkport, Maine, 1994.
4. M Elhadad and K R McKeown. Generating connectives. In *COLING-90*, pages 97–101, 1990.
5. G Gazdar, E Klein, G Pullum, and I Sag. *Generalized Phrase Structure Grammar*. Blackwell, Oxford, 1985.

6. H P Grice. Logic and conversation. In P Cole and J L Morgan, editors, *Syntax and Semantics*, volume 3. Academic Press, New York, 1975.
7. B J Grosz, A K Joshi, and S Weinstein. Providing a unified account of definite noun phrases in discourse. In *Proceedings of the 21st ACL Conference*, pages 44–50, Cambridge, MA, 1983.
8. B J Grosz and C L Sidner. Attention, intentions, and the structure of discourse. *Computational Linguistics*, pages 175–203, 1986.
9. 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.
10. E Hovy. Automated discourse generation using discourse structure relations. *Artificial Intelligence*, 63:341–385, 1993.
11. E H Hovy. Planning coherent multisentential text. In *Proceedings of the 26th ACL Conference*, pages 163–169, Buffalo, 1988.
12. 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.
13. 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.
14. E H Hovy and K McCoy. Focusing your RST: A step toward generating coherent multisentential text. In *Proceedings of the 11th Cognitive Science Conference*, pages 667–674, Ann Arbor, 1989.
15. A Knott. New strategies and constraints in RST-based text planning. Master’s thesis, Department of Artificial Intelligence, University of Edinburgh, 1991.
16. A Knott. An empirical methodology for determining a set of coherence relations. Technical Report (in preparation), Department of Artificial Intelligence, University of Edinburgh, 1994.
17. A Knott and R Dale. Using linguistic phenomena to motivate a set of coherence relations. *Discourse Processes*, 18(1):35–62, 1994. Also available as Technical Report RP-34, Human Communication Research Centre, University of Edinburgh, 1992.
18. L Kosseim and G Lapalme. Content and rhetorical status selection in instructional texts. In *Proceedings of the 7th International Workshop on Natural Language Generation*, pages 53–60, Kennebunkport, Maine, 1994.
19. W C Mann. Discourse structures for text generation. In *Proceedings of the 10th International Conference on Computational Linguistics and the 22nd Annual Meeting of the ACL*, pages 367–375, Stanford, 1984.
20. W C Mann and S A Thompson. Rhetorical structure theory: A theory of text organisation. *Text*, 8(3):243–281, 1988. Also available as Tech Report RR-87-190, USC Information Sciences Institute, Marina del Rey, CA.
21. B Meyer and R Freedle. Effects of discourse type on recall. *American Educational Research Journal*, 21:121–143, 1984.
22. J Moore and M Pollack. A problem for RST: The need for multi-level discourse analysis. *Computational Linguistics*, 18:537–544, 1992.
23. J D Moore. *A Reactive Approach to Explanation in Expert and Advice-Giving Systems*. PhD thesis, UCLA, 1989.
24. 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.

25. J D Moore and W R Swartout. A reactive approach to explanation: Taking the user's feedback into account. In C L Paris, W R Swartout, and W C Mann, editors, *Natural Language Generation in Artificial Intelligence and Computational Linguistics*. Kluwer Academic Publishers, 1991.
26. C L Paris. Generation and explanation: Building an explanation facility for the explainable expert systems framework. In C L Paris, W R Swartout, and W C Mann, editors, *Natural Language Generation in Artificial Intelligence and Computational Linguistics*. Kluwer Academic Publishers, 1991.
27. 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.
28. D Rösner and M Stede. TECHDOC: A system for the automatic production of multilingual technical documents. In *Proceedings of KONVENS '92*. Springer-Verlag, 1992.
29. 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.
30. T J M Sanders, W P M Spooren, and L G M Noordman. Coherence relations in a cognitive theory of discourse representation. *Cognitive Linguistics*, 4–2:93–133, 1993.
31. 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.
32. C L Sidner. What the speaker means: The recognition of speakers' plans in discourse. *International Journal of Computers and Mathematics, Special Issue in Computational Linguistics 9 (1)*, pages 71–82, 1983.
33. 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.
34. 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.

This article was processed using the \LaTeX macro package with LLNCS style

| | | |
|--|--|--|
| Jim summoned his nerve and made a break for the door; | $\left\{ \begin{array}{l} \textit{at that,} \\ \checkmark \textit{instantly,} \\ \checkmark \textit{and} \\ \checkmark \textit{after this,} \\ \checkmark \textit{then} \\ \checkmark \textit{as a result,} \\ \checkmark \textit{consequently,} \\ \checkmark \textit{so} \\ * \textit{in conclusion,} \\ * \textit{it follows that} \\ * \textit{to this end,} \\ * \textit{finally,} \\ * \textit{thereby,} \end{array} \right\}$ | pandemonium broke out in the bar. |
| Sue left the country before the year was up; | $\left\{ \begin{array}{l} \textit{in so doing,} \\ \checkmark \textit{thereby} \rightarrow \\ \checkmark \textit{immediately} \rightarrow \\ \checkmark \textit{and} \\ \checkmark \textit{so} \\ * \textit{in short,} \\ * \textit{at that,} \\ * \textit{to this end,} \end{array} \right\}$ | she lost her right to permanent residence. |
| We waited outside Mullen's door for three hours: | $\left\{ \begin{array}{l} \textit{at last,} \\ \checkmark \textit{finally,} \\ * \textit{instantly,} \\ * \textit{at that,} \\ * \textit{in doing this,} \end{array} \right\}$ | he agreed to see us. |
| The footprints are deep, and clearly defined. | $\left\{ \begin{array}{l} \textit{It follows that} \\ \checkmark \textit{Plainly,} \\ \checkmark \textit{That is to say,} \\ \checkmark \textit{So} \\ \checkmark \textit{Therefore,} \\ * \textit{As a result,} \\ * \textit{Finally,} \end{array} \right\}$ | the thief was a heavy man. |
| The number is divisible by four; | $\left\{ \begin{array}{l} \textit{it follows that} \\ \checkmark \textit{as a result,} \\ \checkmark \textit{hence} \\ \checkmark \textit{thus} \\ * \textit{instantly,} \\ * \textit{this way,} \end{array} \right\}$ | it is divisible by two. |
| One night, Van Diemen nailed a pair of clogs to Mortimer's door. | $\left\{ \begin{array}{l} \textit{From then on,} \\ \checkmark \textit{Thereafter,} \\ \checkmark \textit{After this,} \\ * \textit{For example,} \\ * \textit{All in all,} \end{array} \right\}$ | Mortimer stopped coming to the Dog and Duck. |
| Mary held her breath, | $\left\{ \begin{array}{l} \textit{until} \\ \checkmark \textit{so that} \\ \checkmark \textit{and} \\ * \textit{when} \\ * \textit{from then on,} \\ * \textit{obviously} \end{array} \right\}$ | she turned bright red. |

Fig. 10. Result Phrases: Examples of Substitutability