Problems in the Application of Rhetorical Structure Theory to Text Generation.

Thesis submitted in partial fulfilment of the
M. Eng. Sc. (Cognitive Science) Degree.

University of Melbourne

June 1994

Nick Nicholas

Colours of rethoryk been to me queynte;
My spirit feeleth noght of swich mateere.

— Geoffrey Chaucer, *The Franklin's Tale* 726–727.

Statement of Authorship.

To the best of my knowledge, the work contained within this thesis is entirely my own, except that which is appropriately attributed to other authors. Nor has this work, either in part or in whole, been submitted as assessment to any other degree or diploma.

Nick Nicholas

June 1994

Acknowledgements.

Την ευγνωμοσύνη μου και το θαυμασμό μου πρέπει να απευθύνω στην επιστάμενή μου, *Δρα. Λέσλι Στέρλιγκ*, που είχε την υπομονή και τη γνώση να με καθοδηγήσει τους περασμένους έξι μήνες, από τους οποίους τους τέσσερεις τους πέρασα να οδηγώ σε κύκλους. Θα ήθελα επίσης να ευχαριστήσω το *Δρα. Ρόμπερτ Ντέιλ*, που μου προμήθευσε αρκετές χρησιμότατες πηγές την τελευταία στιγμή.[1]

ghotpu'vamvaD tlho' vIja' je: vIyo'na watSonvaD, jup QaQ Damo' 'ej "ghItlhlIj yIruch" reH jIHvaD jatlhmo'; <melbIn> DuSaQ'a' DIvI' nawlogh loHvaD, DIvI'ngan QaQ Damo' 'ej "ghItlhlIj yIruch" jIHvaD lujatlhQo'mo'; moHamaD maHDIrajIvaD, taHmo' ghaH, yaH wIwavtaHvIS qaStaHvIS DIS bID, 'ej rIQchugh, vaj QIH ru' neH SIQpu'; jon Hajeq QeDpInvaD, jIHvaD Qu' nobmo' ghaH, 'ej pagh lup buD vIghaj 'e' HeQmoHmo'; tlhIngan HolloHvaD, tlhIngan De'wI' QonoS je jeSwI' HochvaD, lupmeyvetlh buD'e' bong vIghajbogh waQmo' chaH (nem tlhIngan HoI vIpojmeH, chay' SoQvaD chen mu'tlheghghom pab 'e' qelbogh nger (ChSChMTlhGhP'KNg) vIlo'jaj); robet <lojbab> leSevalIyervaD, lojban De'wI' QonoS jeSwI' HochvaD je, jIHvaD HolQeD lulIHmo' qaSpa' latlh Hoch (jonwI' lI'be' vImojlaHbej, vaj Do' vImojbe'pu'); 'ej 'IntIha kurtlhImvaD, Qaw'wI''a'qoqvaD, loHvaD po'meH wIch Dachu'mo' 'ej Sengmeyvo' muleHmo'. HochvaD Qapla' van je![2]

---

[1]*Appreciation and gratitude must go to my supervisor,* Dr Lesley Stirling, *for having the patience and the know-how to steer me through these past six months, of which I spent four running myself over. I would also like to thank* Dr Robert Dale, *who provided me with several useful references at the last minute.*

[2]*Thanks also to* Fiona Watson, *for being a good mate and telling me to do my thesis; to the* Executive Committee of the Melbourne University Star Trek Club, *for being good Trekkies and not telling me to do my thesis; to* Mohammad Mahdiraji, *for enduring a semester of sharing an office with me while undergoing little if any permanent damage; to* Dr John Hajek, *for employing me and making sure I should never have an idle moment; to all at the* Klingon Language Institute *and the* Klingon Mailing List, *for filling up those idle hours I somehow ended up having anyway (I hope I'll do a Rhetorical Structure Theory (RST) analysis of Klingon one day!); to* Bob 'Lojbab' LeChevalier *and all those at the* Lojban Mailing List, *for getting me into linguistics in the first place (I'd have been pathetic as an engineer anyway); and to* Indra 'The Destroyer' Kurzeme, *for being an all-round legend of administrative dexterity and keeping me out of trouble. Success and salutations to you all!*

Contents.

1. Introduction.

This thesis examines the problems arising out of the convergence of two fields of linguistic scholarship.

The first field is that of text generation by computer. Long considered a 'poor relation' of natural language understanding within computational linguistics, text generation has made significant progress in the past ten years—to the point where advanced linguistic features of text (such as illocutionary force and stylistic diversity) can be exploited by a computer.

The second field is that of text linguistics—specifically, Rhetorical Structure Theory (RST) (Mann & Thompson (1987)—henceforth M&T), a framework developed to account for text structure above the clause level, by hierarchically positing relations between spans of text. Increasingly, RST is being used not only as a tool for analysing the structure of natural language text, but also as a planning aid in text generation. In particular, it is used to help the computer decide on means to linguistically realise these intratextual relations. The most obvious of these means is the use of text connectives like *because* and *if*. But researchers such as Scott & de Souza (1990) (henceforth S&dS) have established that syntactic means for expressing relations between text spans (such as phrasal embedding) can also be exploited by a text generator.

The application of RST to text generation has highlighted a number of problems with the theory; these problems had not been as disruptive while RST was limited to the descriptive domain. I believe these problems affect not only the analytical use of RST, but its computational use as well. In this chapter I set the background for discussing these problems: I review RST itself, computational linguistics, the use of connectives to signal rhetorical relations, and Scott & de Souza's research.

In the remainder of the thesis, I investigate the actual problems I believe arise from applying RST to text generation. In Chapters 2 and 3, I investigate two related problems for the theoretical framework of RST. In Chapter 2, I consider the ontology of RST: what types of relations between text spans should be considered as part of the theory, and between what linguistic entities should the theory posit such relations. In Chapter 3, I consider the possibilities for a systematic taxonomy of RST relations; I evaluate existing schemes classifying relations of the type used by RST, before proposing my own taxonomy. I then use this taxonomy to account for the distribution of certain text connectives.

In Chapter 4, I evaluate Scott & de Souza's proposed heuristics for text generation using RST. I have singled out their study because (more than any of their colleagues) they make several bold claims about how and why rhetorical relations should be marked textually. In addressing these claims, I aim to shed some light on the assumptions often made by workers in text generation, and the extent to which they are justified.

I also attempt to extend Scott & de Souza's text generation heuristics to subordinate clauses; S&dS assume their heuristics still hold in that domain, without explicitly addressing it in their research. In establishing that their heuristics do not uniformly apply to subordinate clauses, I attempt to clarify the particular features of the rhetorical relations involved leading to this inconsistency.

Particularly in Chapters 2 and 3, and to a lesser extent in Chapter 4, my work is primarily a theoretical-linguistic discussion with computational ramifications, rather than vice versa. However, since RST is consistently evaluated in this work according to how it is applicable to text generation, most of the literature cited originates in computational linguistics, rather than text linguistics.

1.1. Rhetorical Structure Theory.

Rhetorical Structure Theory (RST) (M&T; Mann & Thompson 1986) is an analytic framework designed to account for text structure in running text above the clause level. It takes clauses[3] as its atoms, and relates them hierarchically, using a number of predefined rhetorical relations. These relations are defined functionally, in terms of what their intended effect on the reader[4] is. Examples of such relations are JUSTIFY, ELABORATION, PURPOSE, ANTITHESIS, and CONDITION. The full definition of these relations consists of constraints on the text spans related (most relations are asymmetrical, with a *nucleus* span differentiated from *satellite* spans); constraints on the combined span; and a description of the relation's expected effect.[5]

For example, the relation JUSTIFY, between a nucleus span N and a satellite span S, is described as having the effect "R[eader]'s readiness to accept W[riter]'s right to present N is increased" (M&T 1987:11). In other words, if a JUSTIFY relation is posited, then the span S is understood to provide justification for the writer's claim in N. The following example illustrates how this definition is applied (see also Fig. 1.1.):

> 1. The next music day is scheduled for July 21 (Saturday), noon–midnight.
> 2. I'll post more details later,
> 3. but this is a good time to reserve the place on your calendar.
>
> In this text, units 2–3 are in a JUSTIFY relation with unit 1. They tell readers why the writer believes he has the right to say unit 1 without giving 'more details', in particular without giving the location of the music day event. (M&T 1987:10)

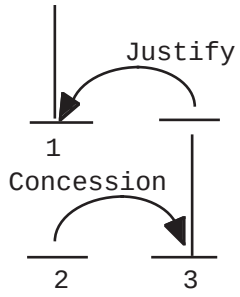

Fig. 1.1. RST analysis of 'Music Day' text.

In ascribing effects to relations between text spans, RST invokes functional linguistic motivation. Each relation posited between any two text spans in RST is described in terms of speaker intent:

> The analyst effectively provides a plausible reason the writer might have had for including each part of the whole text. (M&T:5) For each relation […] definition, the definition applies only if it is plausible to the analyst that the writer wanted to use the spanned portion of the text to achieve the Effect. (M&T:19)

This does not mean that all rhetorical relations are differentiated according to interpersonal meaning alone (to use the systemic-functional term[6].) Rather, according to their

---

[3]I will go into more detail on what constitute the atoms of RST in Section 1.5 and Chapter 2.

[4]Because they primarily analysed written texts, M&T consistently refer to readers and writers, rather than listeners and speakers. I use the same terminology in this work.

[5]See Appendix A for a complete list of definitions of the rhetorical relations presented in M&T.

[6]M. A. K. Halliday's Systemic-functional linguistics (Eggins 1994) distinguishes between three types of meaning. *Ideational* meaning, which represents reality in the world (and can be encapsulated by formal

(continued overleaf….)

Effect, rhetorical relations may be subdivided into two types. *Subject Matter* or *Informational* (Semantic) relations are intended to make the reader *recognise* that there is an Ideational (real-world–describing) meaning relation between the two text spans. Although such relations may conceivably have other perlocutionary effects in context, this recognition is the only perlocution they are *defined* as conveying. Because their perlocutionary effect is so straightforward, these relations are pragmatically uncomplicated, and can be readily analysed by a truth-conditional semantics. Examples of such relations are ELABORATION, CIRCUMSTANCE, PURPOSE, CONDITION, and SUMMARY.

Presentational (Pragmatic) relations, on the other hand, are intended "to increase some *inclination* in the reader" (M&T:18; my emphasis.) This means that they have a non-trivial perlocutionary effect, not limited to mere reader recognition. For example, JUSTIFY has the effect of increasing the reader's inclination to accept that the writer is entitled to her[7] assertion. An informational relation like CIRCUMSTANCE has no such interpersonal effect; it does not attempt to make the reader do anything but accept the model of the world the text is describing. Examples of presentational relations are ANTITHESIS, JUSTIFY, CONCESSION, and EVIDENCE.

It can be argued that most, if not all discourse produced by a language user has a perlocutionary effect. As discussed, Informational relations have a perlocutionary effect as much as Presentational relations do. The distinction is that Informational perlocutions are rather straightforward (getting the reader to recognise some fact about the world), whereas Presentational relations have a much more varied perlocutionary repertoire.

Presentational relations are thus of particular interest from a Speech Act Theory perspective, particularly since there is a one-to-one mapping between the rhetorical relation and the intentionality of the text. EVIDENCE serves to argue a point; MOTIVATION, to encourage a course of action; CONCESSION, to convince that a statement holds despite a seeming incompatibility with its environment; and so on. In fact, the theory is called 'Rhetorical' precisely because it deals with these types of illocutions.

On the other hand, as Moore & Pollack (1992) have established, there is no one-to-one relation between intention and rhetorical relation for Informational relations. Given a Presentational relation, we have a good idea why W wrote the text—what she hoped to achieve by it. Given an informational relation, we know that W intended to *inform* R; but we can only guess what W's underlying motivation was. There are many possible intentions W may have had; to select between them, R has to use a model of W. Moore and Paris (1993) give the example of a computer interlocutor intending to tell a user which screwdriver to pick up. This is a single intention; whereas it would be realised by a single presentational relation (were one available), it can be realised by any number of informational relations: CIRCUMSTANCE *(It's in the top drawer of the toolbox);* ELABORATION *(It's the yellow one);* CONTRAST *(Not the screwdriver you used before; the other one),* and so forth.

In contrast to other accounts of discourse coherence, RST uses an explicit, hierarchical tree structure to describe discourse structure. RST analyses are thus analogous to analyses in traditional syntax. Its use of rhetorical relations hearkens back to researchers such as Grimes (1975), van Dijk (1977) and Longacre (1983). Yet a large subset of these relations has a pedigree of over 25 centuries, ultimately originating in the traditional analysis of persuasive speech, rhetoric. This means that the metalanguage of RST, as well as its representational language, is widely known—another point in its favour.

---

semantics of various sorts); *Interpersonal* meaning, which covers the interactions and affects of agents involved in discourse (covered by accounts such as Speech Act Theory); and *Textual* meaning, which deals with the organisation and presentation of text, as captured by discourse semantics.

[7]For a gender-neutral anaphor, I use either *she*, or the increasingly prevalent singular *they*.

RST has been used widely in the discourse analysis community. For example, Fox (1987) uses it to give an account for the distribution of anaphora in expository English texts. As an analytic tool, it has proven flexible and powerful, and seems to have gained a certain pre-eminence in the field. However, it is an empirical theory, without secure theoretical underpinnings. I believe the application of RST to text generation has highlighted how insecure these underpinnings are, and this motivates the discussion in Chapters 2 and 3.

1.2. Computational Linguistics.

Much of the impetus for RST, and for similar work on text linguistics (such as Grosz and Sidner's Discourse Theory (Grosz & Sidner 1986)) comes from computational linguistics. Mann & Thompson speak of having developed RST

> in the context of work on text generation, designing computer programs that have some of the capabilities of authors; RST thus has both analytical and constructive uses. (M&T:2)

A theory providing a generative account of text structure, conceived explicitly in terms of rhetorical persuasion, has obvious advantages—particularly for text generation, which typically involves instructional text.[8] Moore and Paris (1993) and Hovy (1991) have argued extensively for the superiority of RST, particularly over McKeown's *schemata* (McKeown 1985), as a text planning tool. One of the major attractions RST is perceived as having is its ability (at least in Presentational relations) to encode the intentionality associated with particular utterances—something which schemata, lacking an internal structure, cannot:

> [Schemata] do not include an explicit representation of the effects that individual components of a schema are intended to have on the hearer, or of how these intentions relate to one another or to the rhetorical structure of the text. This presents a serious problem for a system that must participate in a dialogue where users can ask follow-up questions […] If a system does not keep a record of the intentions behind its utterances, it cannot determine what went wrong when the user indicates that an explanation was not completely understood, nor provide an alternative explanation to correct the problem (Moore & Paris 1993:656)

Of course, in its perlocutionary account of rhetorical relations, RST is no more powerful than an approach based directly on Speech Act Theory would be. However, while there have been approaches to text generation based explicitly on Speech Acts (notably Appelt (1985)), Moore & Paris (1993:654) conclude that such approaches leave the linguistic organisation of texts underspecified (particularly for rhetorically complex text spans), and are not computationally feasible, since they require an overly detailed model of the human interlocutor's beliefs.

RST's ability to reflect intentionality is investigated thoroughly in Moore and Paris' work, and is most immediately applicable to interactive discourse—though Vander Linden *et al.* (1992:184) have claimed RST's suitability here is "questionable". But RST is also useful in non-interactive contexts: an awareness of intentionality is central to text planning as a whole, not just error recovery. The linguistic structure of human discourse reflects humans' intentionality (the very *raison d'être* of Speech Act Theory.) What humans plan on saying and how they say it is contingent on what they intend; and computers using natural language need to emulate human ways of constructing discourse. So computers need to emulate human text planning strategies, which can only be done with a model of

---

[8]For example, Rösner and Stede (1992) and Stede (1992) concentrate on the automatic generation of car repair manuals; Moore and Paris (1993) on an advisory system for computer novices; Vander Linden *et al.* (1992) on the automatic generation of telephone manuals. The other main trend in text generation is information retrieval from databases; de Souza *et al.* (1989) use a criminology database, while McKeown (1985) and Hovy (1990) use a naval database.

intentionality. Inasmuch as RST can capture intentionality, it too is appropriate for text planning.

The hierarchical nature of RST also allows flexibility in fleshing out text content as required. For example, Hovy's (1991) operationalisation of RST exploits *growth points*—points on the RST tree which allow the interpolation of supplementary rhetorical information (such as BACKGROUND or ELABORATION) at structurally appropriate places in the text. A block-based system without internal structure, like McKeown's schemata, would not allow such interpolation; this diminishes the system's explanatory flexibility.

RST is becoming an increasingly prevalent tool in text generation planning (de Souza *et al.* 1989, S&dS, Hovy 1990, Hovy 1991, Stede 1992, Rösner & Stede 1992, Vander Linden *et al.* 1992, Hovy 1993, Maier & Hovy 1993, Krifka-Dobeš & Novak 1993, Vander Linden 1993, Moore & Paris 1993). Although much computational RST work is associated with the Information Sciences Institute of the University of Southern California (where Mann, Thompson, Hovy and Paris all work), it is currently used throughout the text generation community.

1.3. Connectives and Rhetorical Structure.

There are a number of textual means to signal that a rhetorical relation holds between two text spans. The most obvious of these are connectives, which can be used to signal most RST relations. Thus, SUMMARY can be signalled by *in all* or *so*; CONTRAST by *but* or *however*; PURPOSE by *to* or *in order to*; CONCESSION by *although*; and so on.

Are connectives *necessary* for a reader to identify that a particular rhetorical relation holds? Mann & Thompson take pains to dissociate their rhetorical analysis from any particulars of linguistic form such as connectives:

> The applicability of a relation definition never depends directly on the form of the text being analysed; the definitions do not cite conjunctions, tense, or particular words. RST structures are, therefore, structures of functions rather than structures of forms." (M&T:19)

Mann & Thompson (1986) make three further contentions about the connective marking of *relational propositions.*[9] Firstly, relational propositions arise independently of any textual cues. In a sequence like *I'm hungry. Let's go to the Fuji Gardens*, there is no explicit textual marker of the relation posited, SOLUTIONHOOD. While some relations may be signalled, it is not *necessary* to do so, they claim, for any rhetorical relation.

Secondly, connectives do not in fact *signal* relations at all; rather, they

> *constrain the interpretation of relational propositions* […] it is the implicit *relations* which are important, with the conjunctions acting occasionally to constrain the range of possible relational propositions which can arise at a given point in a text." (Mann & Thompson 1986:71)

Lastly, even as constraints, connectives operate 'loosely', in that a relational proposition can correspond to many connectives (for example, CAUSE can be associated with *so*,

---

[9] A *relational proposition* is an implicit proposition, arising from a text, that two parts of that text are related in a given way. Mann & Thompson make it clear that relational propositions and rhetorical (RST) relations are equivalent:

> "The relational propositions correspond to the relations of the RST structure of the text. One relational proposition arises from each relation of the text… Recognizing the relations of a text, which is tantamount to recognizing its RST structure and the basis of its coherence, is thus essential to understanding the text." (M&T:20)

For this reason, any conclusions involving relational propositions in Mann & Thompson (1986) should apply to RST relations as well.

therefore, *consequently*, *thus*, etc.), and a connective can correspond to many relations (thus, *but* can be associated with both CONCESSION and ANTITHESIS.) In other words, the mapping between form and function for connectives is not one-to-one, but many-to-many.

Both Bishop (1993) and Blakemore (1989) would dispute the claim that *but* is never a necessary signal for its relations (although Bishop concedes that the claim that connectives merely constrain "does seem generally valid.") In Bishop's example, the deletion of *but* in the following text makes it "not at all clear what connection the reader is expected to make:"

> 3. The cuisine remains the envy of local hotels.
> 4. A warm but unassuming hospitality still awaits.
> 5. [But] now, guests can enjoy the results of extensive renovations.
> 6. Including the newly refurbished Le Club President…
>
> […] *Now* certainly signals a recent change, but there is a sense in which the issue of potential incompatibility is side-stepped, and lines 3–6, without *but*, tend to read like a simple list of the hotel's positive attributes. That is, the 'concessive' sense is lost. It may be, therefore, that certain rhetorical relations do in fact generally require an overt signal, in order to be differentiated from a simple JOINT[10] relation. (Bishop 1993:81)

Blakemore gives a similar argument (1989:27). She mentions the failure of *John is a Republican, [but] he's honest* to generate the implicature *Republicans are not honest* when the connective *but* is omitted ("the hearer might never have accessed the contextual assumption(s) necessary for the derivation of [this implicature]".)

Blakemore works in the framework of Relevance Theory. Therefore, she is not concerned with the relation between text spans, as much as the constraints on implicature following from their juxtaposition. However, the mechanism she describes, impeding the generation of implicatures by denying the 'expectation' generated by the first clause, is the same mechanism of implied contrast which generates the relational proposition of ANTITHESIS, defined as:

> comprehending S and the incompatibility between the situations presented in S and N increases R's positive regard for the situation presented in N. (M&T:12)

So while Mann & Thompson's (1986) claim that rhetorical relations do not *inhere* in connectives is valid, their claim that connectives serve merely to *constrain* rhetorical interpretation is not completely accurate. In the case of *but*, the connective does not merely constrain interpretation to contrastive relations; it makes a previously inaccessible contrastive interpretation (a CONTRAST or ANTITHESIS relation) recoverable for the reader. In that respect, at least, connectives are 'necessary', guiding the reader to make rhetorical sense of the text.[11]

But even where connectives do merely constrain rhetorical interpretation, the fact they do not (always) *signal* particular relations means that the issue of ambiguity remains. It is not always possible to identify which rhetorical relation holds between two spans, on the basis of linguistic evidence. Indeed, it may not even be possible to do so based on contextual evidence.

---

[10] See Appendix A for definitions of all RST relations.

[11] In the context of text generation, of course, the issue of the 'necessity' of a connective is critical, whereas in analytic text linguistics, as shown by Mann & Thompson's position, it is peripheral. This is a good illustration of a recurring theme in this thesis: the inevitable discrepancies between RST as a text-analytic tool, and RST as a text-planning tool.

This kind of ambiguity does not cause any concern for M&T, who write that "ambiguity is also normal for RST. Nothing in the definitions of RST constrains it to single analyses." (1987:28) But they do not imply that there is any pattern to the ambiguity, or any way of predicting what kind of relations the interpretation will be constrained to. In other words, they do not have a principled account of ambiguity in RST—and I feel this diminishes the framework's explanatory power, and the reproducibility of its analyses.

I believe Mann & Thompson are here ignoring a major reason for the practical difficulty in determining rhetorical relations: rhetorical relations do not exist in isolation from each other; they fall in classes. This means that the many-to-many relations between connectives and relations are not completely 'loose': they may not always specify a single relation, but they will typically constrain interpretation to an *interrelated* subset of relations. Failure to acknowledge the family structure of RST relations is, I believe, the cause for a lot of confusion, and I will return to this issue throughout the thesis—particularly in Chapter 3.

1.4. Scott & de Souza's Hypothesis.

Having considered whether connectives are *necessary* to signal rhetorical relations, I now consider whether they are *sufficient*.

I intend to do this within the framework set by S&dS. In this article, the authors appeal to work on psycholinguistics and text understanding, in order to set three requirements for computer-generated texts to achieve "clarity and conciseness in the textual expression of the message at the rhetorical level":

> First, they must be sensitive to the communicative context in which they are set, i.e., one where the writer is an artificial interlocutor, with few resources for predicting or judging the impact of a text on the reader. Second, the chosen expression of the message must be a valid and unambiguous textual rendition of its rhetorical structure (i.e. the rhetorical structure of the message must be derivable from the text.) Third, the chosen expression must be the most easily processed member of the set of all valid and unambiguous expressions of the message. (S&dS:48–49)

They justify the desire for explicit, recoverable rhetorical structure by arguing that

> [s]ince artificial interlocutors clearly have fewer possibilities to make reliable assessments of their audience's ability to 'get the message' than do their human equivalents, their expressions of the message often need to be more explicit than would be ideal. This is particularly the case with respect to the rhetorical aspects of messages, whose understanding generally relies heavily on common-sense knowledge. (S&dS:49)

They go on to claim that, according to their examination of British English and Brazilian Portugese, all rhetorical relations proposed in M&T can be marked textually—whether lexically, phrasally, or syntactically.

Scott & de Souza thus claim, not only that all M&T's rhetorical relations can be marked textually, but that they should be marked *unambiguously*. These are bold claims, which clearly deliniate the different approaches S&dS and M&T take to RST. S&dS's requirement of unambiguous signalling is alien to M&T's view of rhetorical relations, discussed above, where connectives merely loosely constrain rhetorical interpretation. But this is a natural consequence of the fact that S&dS are not concerned with natural language texts, but with computer generation of texts. As they argue, the pragmatic impoverishment of computer texts means they should be subject to evaluative criteria different to those used in text analysis.

But both M&T and S&dS are concerned with the one set of linguistic resources made available by the language; they do not each have access to a different body of connectives. And the same connectives are bound to carry the same kinds of rhetorical ambiguity. So it

should not come as a surprise that S&dS's claim (at least in its strong version) is not substantiated by their empirical work. S&dS refer the reader to de Souza *et al.* (1989), but in that article the authors explicitly state that the LETTERA project uses only a subset of RST relations (de Souza *et al.* 1989:232), and they allow the use of the notoriously ambiguous *and* in CAUSE as well as SEQUENCE relations (de Souza *et al.* 1989:229).

In fact, even in their 1990 article, Scott & de Souza find they cannot unambiguously mark every relation defined in M&T. When confronted with five closely related relations indicating causality (VOLITIONAL CAUSE, VOLITIONAL RESULT, NON-VOLITIONAL CAUSE, NON-VOLITIONAL RESULT and PURPOSE), all of which are marked by *because*, they hedge out of the problem by pointing to the CAUSE *cluster* of relations, defined in M&T to include all these relations. S&dS go as far as to add EVIDENCE to the cluster (S&dS:52), allowing *because* as a lexical signal for any member of the cluster.

M&T themselves distinguish between CAUSE and RESULT as a matter of nuclearity: the satellite clause in these relations denotes *result* and *cause*, respectively. Since it has a different *syntactic* function in these two types of relation, *because* is not ambiguous in distinguishing between them: the underlying relational predicate is the same. The volitionality distinction, on the other hand, is semantically motivated: it distinguishes between the predicates VOLITIONALLY-CAUSES and NON-VOLITIONALLY-CAUSES in the relational proposition. I discuss whether this distinction is relevant to a rhetorical theory in Chapter 2.

As I will argue, in speaking of a CAUSE cluster, M&T are making the first steps towards recognising an RST taxonomy. S&dS pick up on this; their extension of the cluster also shows awareness of such a taxonomy. But I believe this taxonomy is ultimately incompatible with their avowed aim of non-ambiguity.

Having accepted M&T's claim that connectives serve to constrain, rather than to mark rhetorical relations—a claim illustrated so well by the connective *and*—S&dS formulate a weaker version of their aim: that this natural ambiguity of connectives should be limited as much as possible in text generation:

> Although some degree of ambiguity will have to be tolerated, ambiguities that arise from the generation of a very weak marker that also happens to be a stronger marker of another rhetorical relation should not count among them. (S&dS:52)

In other words, if at all possible, use a connective that will mislead the reader as little as possible—given (so the assumption goes) that we *can* identify less misleading ('prototypical') textual markers for a given relation. Unfortunately, rhetorical marking tends to be a subtle business, in which any notion of connective non-ambiguity will be overridden by many conflicting factors. Stede (1992) is an example of a quite realistic approach to rhetorical marking, where the stylistic constraints under which a text generator works are described much more explicitly. More often than not, these kinds of constraints radically affect the choice of connective to be used in a text. So it will often prove impractical to dedicate a given connective to the task of signaling a given relation in generating text.

In their work, S&dS actually ignore stylistic constraints as being distinct from intelligibility constraints:

> The ultimate criterion of what it means for a text to be good is thus a cognitive rather than a strictly linguistic one: the easier it is for the reader to decode the intended message from the text, the better the text will be. (S&dS:47)

The current feeling amongst workers in text generation (whose linguistic sophistication has increased significantly since the mid–'80s) is that 'strictly linguistic' criteria (such as repetitiveness, or genre constraints on lexicon and syntax) matter a great deal to the ac-

ceptability of text output. However, this does not affect any examination of how valid S&dS's main contentions are.

S&dS have set up a system of heuristics to exploit a plan for text generation formulated in RST. These heuristics are intended to make the generated text's rhetorical structure more explicit, and thereby to make the text easier to process and to understand. This is a sensible idea, and the authors have been careful to cast it as a set of heuristics, rather than anything as iron-cast as an algorithm. Nonetheless, not all the consequences of applying this heuristic have been fleshed out in S&dS. I believe there are issues latent in its applicability, which are relevant not only to text generation, but also to the investigation of connectives in natural text, as signals of rhetorical relations.

1.5. The issues considered.

In this thesis, I investigate three areas where the application of RST to text generation has given rise to difficulties. I am writing from a linguistic, rather than a purely computational perspective; for this reason, I do not consider implementation issues. This work investigates the linguistic ramifications of these difficulties, rather than explicitly proposing a computational model to resolve them. Nevertheless, the critiques I make should prove relevant to improving planning and output in text generation.

The areas I consider are as follows:

1.5.1. RST Ontology (Chapter 2).

There are discrepancies between RST used as an analytic tool in descriptive linguistics, and RST as a planning tool in text generation. The most notable of these is the issue of what the scope of RST should be. The smallest items Classical[12] RST ever considered were clauses; there was no attempt made to analyse intraclausal rhetorical relations. In text planning, on the other hand, the set of atoms to be organised by RST is not a set of found clauses, but a set of unintegrated *propositions* in the text generator's knowledge representation language—which need not all be realised as distinct clauses.

As a result, the RST trees drawn by text generation researchers consider nominalisations, relative clauses, and even adjectives as satellite text spans, to be linked to nuclei within the main clause. Classical RST would consider these text spans as embedded within their matrix clauses, and thus not subject to rhetorical analysis. I give examples of this discrepancy, and of the ontological questions this gives rise to for RST.

A second major problem, which has become obvious as recent computational work in RST has arbitrarily expanded its inventory of relations, is what types of relations a rhetorical theory should be describing. This becomes even more relevant as RST starts to encompass intraphrasal relations. What, if any, should the boundary be between truth-conditional semantics and rhetorical analysis? What should a rhetorical theory do with a text like *Heat the pot until it starts frothing?* There is obviously a coherence relation between two clauses in this text, since the interpretation of the anaphor *it* depends on the earlier *the pot.*[13] But how should the coherence relation be specified? Is it merely a CIRCUMSTANCE, as in M&T? Should it be specified as a temporal relation, as done by Hobbs (1985)? Should it be specified even further, as a punctiliniar temporal relation, as in Halliday & Hasan (1975)? At this level of semantic analysis, what distinguishes RST from tense logic, or propositional semantics in general? And since some RST relations are

---

[12] I use the term 'Classical', after Rösner & Stede (1992), to refer to M&T's original formulation of RST, as distinct from its subsequent modifications by various text generation researchers.

[13] Although *it* and *the pot* are not co-referential; a bridging inference is required to enable the anaphor to be resolved, and thus, for coherence to be established.

necessarily not propositional-semantic, where does the boundary between rhetorical theory and propositional semantics lie?

In considering these issues, I am addressing a widespread intuition amongst researchers in text linguistics, that certain types of relation 'belong' in a rhetorical theory, and others don't. I attempt to provide an formal, explicit formulation of this intuition. Although my results are not conclusive, I believe they will help delimit RST ontology more clearly than has been done in the past.

1.5.2. RST Taxonomy (Chapter 3).

A disturbing property of RST, used both descriptively and in text planning, is the tendency of analysts to treat its set of relations as open-ended, and to expand it arbitrarily. This open-endedness has been part of the theory since its inception, and there are many instances of it in the literature. For example, Bishop (1993) adds CONSISTENCY, CORRECTION and DISJUNCTION relations to the RST inventory; Rösner & Stede (1992) add PRECONDITION, UNTIL, ALTERNATIVE and STEP–SEQUENCE. Mann & Thompson themselves do not provide any external justification for their inventory of rhetorical relations. While a small set of relations is identified as recurring in text, these relations were arrived at empirically, without any formal underpinnings; and there are several 'twilight-zone' relations, like PRESENTATIONAL SEQUENCE and DISJUNCTION, which M&T considered but did not eventually incorporate into their scheme.

The types of rhetorical relations encountered in a text vary according to that text's genre. So analysts have good reason for wishing to fine-tune their inventory, to better deal with the texts they are working on. But this causes an obvious methodological problem. If each analyst is allowed to choose their own set of rhetorical relations, RST analyses become no longer reproducible. Short of formalising some operational interlanguage, into which the definitions of each analyst's rhetorical inventory can be translated, we will be unable to compare analyses using different inventories. Furthermore, any linguistic generalisations made using one inventory may not be applicable to texts analysed using another. This makes the discrepancy a more serious problem than the structural ambiguities that have become acceptable in RST.

It is not as obvious that this variability is such a problem in text generation. Indeed, Vander Linden *et al.* (1992) welcome it:

> [RST] has been successfully applied to a variety of written genres […]. Generally a new domain has dictated modification to the inventory of relations, but this very adaptability is one of its more useful features. (Vander Linden *et al.* 1992:184)

True, the *analytical* problems caused by a proliferation of rhetorical relations do not get in the way of a text planner's work. And flexibility in applying particular rhetorical relations during planning may be desirable: the set of relations a planner is aware of may have to be augmented, to reflect better the planner's intentionality structures, and to accommodate genre and implementation demands.

But I do not believe the way to make the inventory more flexible is by making it completely open-ended. Appending ever more atoms to an unstructured, indefinitely long list does not achieve the kind of flexibility text generation researchers want. The list can balloon out to an unmanageable size; crucial generalisations may be missed; RST may start being pushed into fields it is not well equipped for, such as tense logic; and the compositionality of RST (its major advantage over rhetorical schemas) may be compromised as text relations become ever more specific, to the point of becoming templates in disguise.

In any case, presumably there are always *principled* reasons for adding relations to the rhetorical inventory. These principles will not be captured, let alone exploited, by making the inventory an unstructured list.

An issue related to these concerns is the ambiguity both of rhetorical relations in text, and of the connectives that signal them. This ambiguity has worried many workers in text generation, who seek to reduce it. Yet there is rhyme and reason to it. Analysts find difficulty in distinguishing certain rhetorical relations, because they intuitively feel these relations have features in common. Likewise, 'ambiguous' connectives do not cover a random subset of the rhetorical inventory; the range of relations they do cover may be explained by identifying features these relations share. And in text generation itself, identifying and emphasising the particular features where the relational options *differ* in an ambiguous text is likelier to produce text less ambiguous than would any approach that treats relations as atomic.

The common thread here is the need for a feature-based taxonomy of RST relations. I examine and critique several proposed rhetorical taxonomies—Longacre (1983), Halliday & Hasan (1976), Hobbs (1985), Mann & Thompson (1986, 1987), Sanders *et al.* (1992), and Maier & Hovy (1993). I then make some suggestions as to what an RST taxonomy should look like; and I use this taxonomy to account for the 'ambiguity' of a few frequent connectives, such as *so* and *and*. This should provide some parameters within which the open-endedness of RST can be circumscribed.

1.5.3. The Scott & de Souza programme (Chapter 4).

As we have seen, S&dS address the issue of rhetorical ambiguity by maintaining that computer-generated text needs to have all its rhetorical relations explicitly and unambiguously signalled by textual marking. Their contention is expressed as a heuristic, and it would be unfair to treat it otherwise. But it is worth asking questions like the following:

- How workable are S&dS's heuristics in a real text generation system?

- Is rhetorical signalling essentially the same at all hierarchical levels of text structure?

- How do connectives and phrasal signals (such as *the reason is* or *because of this)* differ in effectiveness in signalling rhetorical relations?

- Are there rhetorical relations which are intrinsically difficult to signal textually? If so, can an RST taxonomy be used to explain why?

- To what extent should text generation exploit the commonalities between rhetorical relations, which generate much of the ambiguity Scott & de Souza descry?

I look at the workability and the implications of Scott & de Souza's contentions from a variety of linguistic angles. I also consider what justification there is for their claims. Does the fact that certain relations are hard to textually signal *matter*, given the appropriate world-knowledge and genre constraints? How much textual signalling is too little—or for that matter, too much? How strong is the psycholinguistic justification for Scott & de Souza's position? What sort of rhetorical marking is in fact 'necessary'—and for what purpose? Finally, I scrutinise the psycholinguistic literature Scott & de Souza have cited, to determine how secure the foundations of their theory are.

2. RST Ontology.

2.1. What is a rhetorical relation?

As mentioned in Chapter 1, the rhetorical inventory of M&T contains two different types of relation. Presentational relations cover perlocutions; these relations originate from the same linguistic domain as Speech Act Theory. Informational relations are not from this domain—or rather, they are involved with this domain only in a trivial way. Instead, they denote the relationships between events either in the real world, or in our mental representation of it. By involving real world events, rather than speech acts, they set up an ideational, rather than an interpersonal semantics—which can be readily analysed by truth-conditional semantics.

The divide between these two types of relation is, in my view, a basic difficulty with RST. The theory has two distinct modes of analysis. The Presentational mode is informed by Speech Act Theory; the Informational mode, by truth-conditional semantics. Indeed, such a split in analytic levels is to be expected of a theory like RST, since the links between text spans can arise for either primarily presentational or informational reasons.[14]

This divide is much more of a problem for RST as a linguistic than as a computational tool. None of the ramifications of this divide discussed below mean that any text generators need to be redesigned. However, I suspect that computational linguists' lack of awareness of this divide leads to a misunderstanding of rhetorical theory, and of what analytical power should be reasonably expected of it.

One problem that arises from RST's split-level analysis of coherence is that the relation between two text spans can operate at both levels in distinct ways, as Moore & Pollack (1992) and Rösner & Stede (1992) point out. A text like

Come home by 5:00. Then we can go to the hardware store before it closes.

can be analysed both as a MOTIVATION, at the presentational level, and as a CONDITION, at the informational level (Moore & Pollack 1992:541). Since pragmatics and truth-conditional semantics are independent disciplines, a pragmatic, presentational analysis of the text cannot be reduced to a truth-conditional analysis, nor vice versa.

Moore & Pollack find the inability of RST to support two levels of analysis in this way "a serious problem for the theory".[15] However, I wish to consider not this, but a related problem, arising from the two levels of analysis. The problem is the following: to classify the different ways text can cohere, RST invokes distinctions between rhetorical relations; these distinctions can be formulated in either formal semantics or Speech Act Theory. To what extent should RST exploit the resources of these two disciplines—both of which are disciplines distinct from RST itself? Should every distinction between coherence relations that can be formulated in formal semantics or Speech Act Theory, be formulated in RST as well? Is a full propositional semantics a component of RST? For that matter, is a full Speech Act Theory a component of RST?

---

[14]Note that systemic-functional linguistics claims that both types of semantics associated with these relations—interpersonal and informational respectively—are simultaneously present in text. But it should still be possible to distinguish which semantics is primarily involved in determining the type of textual link.

[15]M&T (1987:20) acknowledge the possibility of simultaneous rhetorical analyses, and consider it acceptable. Unlike Moore & Pollack, however, they did not point out that a major reason for simultaneous analyses are the two different levels of analysis RST exploits. Nor do they provide any systematic account for this type of simultaneity, other than recognising it exists.

This is an abstract formulation of the problem; consider the following concrete example. Rösner & Stede (1992) propose several extensions to the RST rhetorical inventory, having applied it to the text generation of instructional manuals. One additional relation they propose is UNTIL, whose effect is defined as "R recognizes that N has to be carried out only as long as S is unrealized", and which is exemplified by *Turn the radiator cap counter-clockwise until it stops*.

More than one rhetorical analyst with whom I have discussed UNTIL found it a peculiar addition to make to the RST inventory. But what is so strange about UNTIL? The relation described by Rösner & Stede does in fact hold between the two text spans in this example. The text spans are not even problematic in any way; this is precisely the kind of hypotaxis Matthiessen & Thompson (1988) believe is a grammaticalisation of Nucleus–Satellite rhetorical relations. And the relation between the text spans is certainly one of coherence.

True, the text is coherent, and UNTIL-hood is as much a relation between two text spans as CAUSE or PURPOSE. But I believe rhetorical analysts find UNTIL objectionable for a very simple reason. To put it informally, they don't believe a rhetorical theory should be in the temporal logic business.

That is to say, the distinction between UNTIL, and, say, a relation like DURING is made (at length) in temporal logic, a formal semantics.[16] If such distinctions were admitted into RST wholesale, there would be a proliferation of different temporal relations. In fact, this already occurs with Halliday & Hasan's (1976) analysis of connectives, which are distinguished according to their semantics as *Punctiliniar, Durative, Repetitive, Interrupted*, etc.

But if a rhetorical theory starts making distinctions between relations at this level, then it becomes as powerful as temporal logic. Taken to its logical conclusion, a rhetorical theory would have to be as powerful as the whole of formal semantics *and* Speech Act Theory, to account for all possible distinctions between its relations. Is this prudent? After all, RST is not designed to have any access to the lexemes on which such a semantics is compositionally based: whatever RST's atoms might be, they are not lexemes. And such a portmanteau of analytical approaches does nothing to make RST any more coherent or modular a theory. It doesn't seem necessary to distinguish between relations at such depth, in order to establish the hierarchical structure of a text.

So does a rhetorical theory need to include a temporal logic component? M&T don't seem to think so, since they collapse all temporal relations in their scheme into CIRCUMSTANCE—a relation which is not even principally temporal. Most workers in discourse analysis would probably intuitively agree with their decision. But is there any objective, explicit way to determine what level of discrimination a rhetorical theory should make? Is there any way of defining the extent to which RST covers the same ground as formal semantics? After all, causation and intention are subject to truth-conditional semantic analysis as much as temporal relations are. Yet few rhetorical analysts would consider CAUSE or PURPOSE superfluous in a rhetorical theory in the same way they would UNTIL.

I believe such an objective criterion for rhetorical discrimination exists, and I now describe it.

---

[16] Admittedly, temporal logic is best equipped to handle non-deontic instances of UNTIL, such as *They turned the radiator cap counter-clockwise until it stopped*. The imperative example given introduces an illocutionary dimension temporal logic is not as well equipped to deal with. I would still argue, however, that the *temporal* aspects of such a directive (as opposed to, say, its conditional aspects) do not seem to fit in a rhetorical theory.

2.1.1. The relational criterion.

In their attempt to establish parameters according to which coherence relations can be classified, Sanders *et al.* (1992) set up what they call 'the relational criterion'. The parameters they consider must satisfy this criterion to be considered valid in distinguishing rhetorical relations:

> A property of a coherence relation satisfies the relational criterion if it concerns the informational surplus that the coherence relation adds to the interpretation of the discourse segments in isolation (i.e., if it is not merely a property concerning the content of the segments themselves.) This does not imply that the meaning of the connected segments is neglected. Because coherence relations connect representations of discourse segments, the meaning of the segments must be compatible with the coherence relation. What the relational criterion does imply, however, is that we will focus on the meaning of the relation and not on the meaning of each specific segment. (Sanders *et al.* 1992:5)

So a feature distinguishing between rhetorical relations is relevant to Sanders *et al.*'s venture (and, I contend, to a rhetorical theory in general), if it involves a difference in the relation between text segments *qua implicated* (defeasible) meaning.[17] This is because the feature is information we *add* to the referential meaning of the individual text spans, by presupposing its corresponding rhetorical relation. If it is added *only* by this presupposition, then presupposing another rhetorical relation will eliminate this feature from the text meaning; so the added information is defeasible.

For example, the fact that the text

> *I'm not going to start learning Dutch. You can't teach an old dog new tricks.*

involves a CAUSE relation can only be conversationally implicated, given our real-world knowledge of learning frames, idiom, conventions for disclaimers—and, of course, the juxtaposition of the sentences. The causal relation does not inhere in the semantics of the individual sentences, but in positing a CAUSE relation between the two sentences. Because of this, the supposition that this is a CAUSE can be defeated by positing a different rhetorical relation. For example, my not learning Dutch could be mentioned as EVIDENCE for the generalisation *you can't teach an old dog new tricks.*[18]

On the other hand, if the feature distinguishing the rhetorical relations does *not* reside in the actual relation between the text segments, but is local to the segments as independent semantic units, then it is not relevant to a rhetorical theory, since it does not involve *relational* meaning. The presence or absence of such a feature is not defeasible: it is directly deducible from segment semantics; and it is not information *added* by positing a rhetorical relation, since it arises independently of any rhetorical theory. Positing a different rhetorical relation in the text, in order to defeat the feature, will merely yield nonsense like *I ate lunch at 1 P.M. ??Later that day, Fred ate lunch at 11 A.M.*[19]

---

[17] If the feature concerns the informational surplus involved in a relation, the feature *may* still be manifest, or overtly signalled textually (for example, causation is overtly signalled by the connective *because*). And any connectives used to signal the relation would still have a literal, entailed meaning; their meaning cannot be defeated by ascribing a different rhetorical relation to the span. There is no way that the EVIDENCE interpretation of *Since there is ash on your lapel, **it follows that** you have been smoking* could be defeated, since such a defeat would clash directly with the entailed semantics of *it follows that.*
But it should still be possible for the feature in question to be merely implicated, and therefore defeasible, *if no textual signal of the feature is present*. For example, *there is ash on your lapel; you have been smoking* could be construed as CAUSE just as easily as EVIDENCE.

[18] This defeat of implicature involves a shift in nuclearity; but nuclearity is relational information, and cannot be the property of a rhetorically isolated text span—which is what the truth-conditional translations involved in applying the relational criterion are. So the relational criterion is not violated.

[19] The phrase *later that day* fails to defeat the BEFORE relation which follows directly from the propositional semantics of the individual sentences. This would count as evidence against Temporality being a feature relevant to a rhetorical theory, according to the relational criterion.

To generalise from this: if the feature distinguishing the rhetorical relations involves the relation *qua* formal-semantically *deduced*, rather than *implicated* meaning, then it is not relevant to a rhetorical theory. (This assumes that such a deduction is contingent on the compositional meaning of the individual text spans, which would not admit any 'informational surplus' outside this compositional meaning.)

As an example of such a rhetorical feature, consider the distinction made by M&T between VOLITIONAL and NON-VOLITIONAL CAUSE. If the Volitionality parameter fails the rhetorical criterion, then texts involving VOLITIONAL CAUSE are still coherent—but the distinction between the two is not appropriate to make in a rhetorical theory: it is not relevant to the rhetorical *relation* between two texts, but, presumably, to the propositional semantics of the related texts in isolation. Hence, it should be analysed within formal semantics instead. In fact, this is what I argue in the following section.

2.1.2. Testing the relational criterion.

A test for whether a rhetorical feature passes the rhetorical criterion can be outlined as follows: Take two rhetorically related text spans. Delete any textual signal of their rhetorical relation, and translate the two into a truth-conditional semantic notation (which would include, for example, time information, but exclude real-world, contingent information such as frames or communicative conventions.) Can the particular type of relation postulated between the two texts be implicated and defeated? Then the relation type is relevant to a rhetorical theory. Does the type of relation follow directly by deduction from the individual texts' semantics? Then it is not relevant.

A problem with this test would seem to be the property of adversative relations pointed out by Bishop (1993) and Blakemore (1992), and described in Chapter 1. In the absence of an explicit textual marker (like *but),* it is often impossible to recognise that a text is meant to convey CONTRAST, rather than JOINT.

But the test as described would not actually conclude that adversativity is rhetorically irrelevant. Implicating a CONTRAST between the 'uncoupled' texts *John is a republican* and *John is honest* is still possible. The 'contextual assumptions' Blakemore believes would be unavailable to a reader, in the absence of *but*, would be made available by the reader presupposing that the texts are in CONTRAST. The reader should then be able to create a mental framework of presuppositions to support this CONTRAST. This framework can subsequently be defeated by positing a new rhetorical relation, and therefore a new mental framework.

The relational criterion offers a bold solution to the problem of where the boundary lies between rhetorical theory and truth-conditional semantics. Any relations that lie within the analytic domain of pragmatics—whether Speech Act Theory proper, or merely conversational implicature theory—are considered rhetorically relevant. Any relations lying outside pragmatics, and therefore directly subject to a truth-conditional semantic analysis, are automatically excluded from rhetorical consideration. In effect, the relational criterion firmly delimits rhetorical theory as a pragmatic theory, to have no overlap with truth-conditional semantics.

I do not think this is a flaw in the criterion. The implicated, defeasible nature of features that pass the criterion is more consistent with Mann & Thompson's (1986) description of relational propositions than are the truth-conditionally semantic features that fail it. Given the pragmatic orientation of rhetorical structure, and the fact that a well-developed propositional semantics already exists independently of rhetorical theory, it seems to make more sense to subdivide RST's Informational categories in a way consistent with pragmatics, instead of immersing the theory into semantic categories foreign to it. That is precisely what the relational criterion achieves.

In the remainder of this section, I apply the relational criterion to coherence relations, as a basis for argumentation on whether particular meaning features belong in a rhetorical theory. The features considered are Volitionality, Temporality, and Elaborativeness. The argumentation for each feature yields a different answer, and also serves to highlight some possible shortcomings of the criterion.

While I believe the criterion gives the right answer for Volitionality, I do not feel as confident with its results for the other two relations. I would certainly not claim the criterion is a definitive answer to the levels-of-rhetorical-analysis problem. Nor would I claim that it fully captures the intuitive reasons why discourse linguists tend to reject rhetorical relations like UNTIL. But I do believe the relational criterion is an essential first step, in a field where no such attempt to delimit rhetorical ontology seems to have been made before.

2.1.3. Against Volitionality.

Showing that Volitionality is irrelevant to a rhetorical theory is independently motivated by the fact that it does not seem to impinge on the interclausal structure of text (the same connectives—*because, so, since,* etc.—are used for volitional and non-volitional causation.)[20] And the decision on whether it applies to a text span (once causation is established) seems to depend more on world-knowledge and inferencing involving potential agents, than inferencing involving potential types of relation.[21]

A proof of that Volitionality is rhetorically irrelevant, involving the relational criterion, would run as follows. Consider the differences in wording between M&T's definitions of VOLITIONAL and NON-VOLITIONAL CAUSE:

relation name: VOLITIONAL CAUSE.	**relation name:** NON-VOLITIONAL CAUSE.
constraints on N: presents a volitional action or else a situation that could have arisen from a volitional action.	**constraints on N:** presents a situation that is not a volitional action.
constraints on the N+S combination: S presents a situation that could have caused the agent of the volitional action in N to perform that action.	**constrains on the N+S combination:** S presents a situation that, by means other than motivating a volitional action, caused the situation presented in N.
example: *I'm going to the corner. The walk will do me good.*	**example:** *I went riding last week. I was sore for three days.*

Table 2.1.: Differences between VOLITIONAL CAUSE and NON-VOLITIONAL CAUSE.

The differences between the constraints on N are, of course, non-relational: they do not alter the *relation* that holds between N and S. The differences between the constraints on N+S *are* relational; but they can be argued to be directly deducible from the non-relational constraint on N. If the constraints on N are swapped (which should be possible if the constraints on N and on N+S are logically independent), then the constraints on N+S no longer make any sense. S cannot, hypothetically, both cause an agent of a volitional action to do something, and actually cause a non-volitional action instead. If it could, we would surely say it did so without actually motivating a volitional action in the first place. Conversely, S cannot cause a volitional action by any means other than motivating a volitional action.

---

[20]The fact that Volitionality or Animacy *is* linguistically realised by certain structures within the clause (even in English) is irrelevant to this point: it does not affect the linguistic realisation of the *relation*.

[21]M&T justify the distinction by claiming that, in defining causal relations, "it is hard to include both situations that are intended outcomes of some action and causation that does not involve intended outcomes, such as physical causation." (M&T:57) But I am not convinced that the disparity in animacy (a formal-semantic category) between the two types of causation should motivate a *rhetorical* distinction between them.

To maintain that these hypothetical situations are possible would be sophistry. Therefore, a volitional nucleus necessitates a volitional causation, and a non-volitional nucleus—a non-volitional causation. But this means that the constraint on N+S follows directly from that on N. The conclusion is that the only material distinction between VOLITIONAL and NON-VOLITIONAL CAUSE is local to the relation nucleus—and thus falls outside the scope of a rhetorical structure theory, according to the relational criterion.

2.1.4. Against Temporals.

To return to the example with which this discussion began: Sanders *et al.* (1992) argue that temporal relations (including UNTIL—and, indeed, SEQUENCE) do not belong in a rhetorical theory as relations distinct from simple additive relations:[22]

> [T]emporal relations belong to the classes of additive relations and […] the properties distinguishing temporal relations from other additive relations concern the referential meaning of the individual segments. […] Given the tense and the aspect of the segments, the temporal properties of two related segments are more or less fixed. A first consequence is that in an unmarked sequence of two segments, the reader does not have the freedom to ignore the temporal meaning aspect. A second consequence is that the order of the segments in a temporal sequence cannot be reversed freely without disturbing the coherence relation. (Sanders *et al.* 1992:28)

In other words, the argument against distinct temporal rhetorical relations runs the same way as the argument against volitionality. Consider an UNTIL relation, like *They turned the radiator cap counter-clockwise until it stopped*. We can take it as given that the nucleus and the satellite are in a CIRCUMSTANCE relation. We can also deduce, purely from the referential meaning of the individual text spans (so the argument goes), that the nucleus event ceases at time t_1, and the satellite event occurs at time t_1. Given all this information, the relation can only be one of UNTIL: no other temporal relation is possible.

In that case, to say the relationship is UNTIL is to say nothing new, since we already know it to be a CIRCUMSTANCE (given rhetorical information), in which one event ceases upon the occurrence of the other (deduced semantic information.) To refine CIRCUMSTANCE into UNTIL, they argue, is no longer the business of a rhetorical theory: positing an UNTIL relation adds nothing to the text that could not have already been derived from propositional semantics.

Unfortunately, this is not strictly true. Given the tense and aspect of *they turned the radiator cap* and *it stopped*, with the connective *until* removed, a BEFORE or WHEN relation could still be implicated between the two; it is only world-knowledge that makes us select UNTIL.

In fact, Sanders *et al.* have phrased their programme rather loosely. For the relational criterion to reject temporal relations, the referential semantics of the individual segments must have access, not to the linguistically overt tense and aspect of the segments (which give inconclusive evidence for discrimination between temporal relations), but to the full temporal setting and event contour denoted by the segment. Sanders *et al.* posit a clause-delimited referential semantics, which includes this denoted temporal information, but excludes relations like causation between segments.[23] But the temporal information is not necessarily linguistically realised; therefore, it is as subject to implicature and defeasibility as causation is.

---

[22]In Sanders *et al.*'s scheme, additive relations include JOINT and CONTRAST—relations that can be expressed with the logical connective *and*, but not with *if*.

[23]Of course, overt causative predicates *realised* within the segment (for example, the predicate in *Shmuckley's theory* ***caused*** *a big stink)* would not be excluded.

So the relational criterion does not give an unequivocal answer to the question of whether temporality is rhetorically relevant. An unequivocal answer would require a more thorough specification of what referential semantics the criterion can appeal to, and how the criterion relates to implicature.

2.1.5. For ELABORATION.

The constraints on the ELABORATION relation, as defined in M&T, are stricter than for other rhetorical relations. For an ELABORATION to hold, an "element of subject matter" in the nucleus must form the focus of the satellite; the satellite and nucleus must be in a part-whole–type relation; and the satellite merely acts to supplement the information in the nucleus, rather than contextualising it (unlike BACKGROUND.)

Because of these stricter constraints, analysts may feel this relation is less 'substantial' than relations like MOTIVATION, CAUSE, or even BACKGROUND: it somehow appears to convey less meaning than these. Furthermore, because ELABORATION satellites are related to their nuclei by virtue of being somehow their embellishments, rather than any more concrete relation, the functional effect seems more local to the satellite. Sanders *et al.* (1992) appeal to this viewpoint to argue that 'descriptive relations'[24] fail the relational criterion:

> [In an ATTRIBUTION,] The 'attributive' meaning can be located in the second segment and therefore is not part of the coherence relation. […] Hence, such descriptive 'relations' are not coherence relations at all. This may be illustrated best by Meyer's (1975) rhetorical predicates that are 'primarily responsible for giving prose its overall organization.' That rhetorical predicates are not identical to coherence relations becomes clear from the fact that rhetorical relations can also relate propositions within a simple clause; see (39), in which the attribute of *having a color* is connected through the rhetorical predicate attribution with the subject *parakeets*.
>
> *(39) parakeets or budgerigars are vividly coloured birds.* (Sanders *et al.* 1992:26)[25]

In light of what I discuss in the following section, it is interesting that Sanders *et al.* mention the fact that Attributive ELABORATION can be realised intraclausally, as an argument against it being a valid rhetorical relation. Vander Linden *et al.* (1992) argue that a nominalisation like *for recharging* is as much a PURPOSE span as a full clause like *to recharge the battery*; the fact that PURPOSE can be realised intraclausally should not mean we have to reject it as a rhetorical relation!

I also don't believe eliminating ELABORATION is consistent with why the rhetorical criterion is desirable in the first place. The purpose of the criterion should be to keep formal semantics from interfering with rhetorical analysis. But the type of relation ELABORATION exemplifies (as I argue in Chapter 3, this type also includes CIRCUMSTANCE and BACKGROUND[26]) is not the kind of semantics we would like to exclude from the theory.

Unlike volitionality or temporality, these principles underlying *these* relations cannot be captured by a referential, truth-conditional semantics. The relationships described by these

---

[24]Sanders *et al.* characterise 'descriptive relations' as including ATTRIBUTION and EQUIVALENT. These relations presumably correspond to M&T's ELABORATION, since M&T consider ATTRIBUTION a subtype of ELABORATION.

[25]To clarify a possible terminological confusion: what I have been calling *rhetorical relations*, following M&T, Sanders *et al.* have been calling *coherence relations*, after Hobbs (1985). These coherence relations include relations such as CAUSE, CONTRAST, EVIDENCE, and so on. Sanders *et al.* are arguing that the (interclausal) relations *they* are considering are not identical to those Meyer has considered, because Meyer admits *intraclausally* realised relations into his scheme.

[26]Interestingly, although Sanders *et al.* cannot fit these two relations into their scheme (as they cut across their Causal/Additive distinction), they agree that *these* relations are important in the functioning of language, and do not dispute that they are coherence relations.

relations are not real-world relations; they involve the organisation and presentation of text. In Hallidayan terms, they involve not ideational, but textual semantics. For that reason, they can only be expressed in terms of discourse analysis. This makes these relational distinctions decidedly relevant to a rhetorical theory, which purports to analyse discourse structure functionally.

In any case, I do not believe that ELABORATION fails the relational criterion the way I have phrased it. This is because the question 'which span is an ELABORATION of which?' cannot be resolved simply by looking at the referential semantics of the individual spans, and because the decision that a text span is an ELABORATION is defeasible.

As an example of how the nuclearity of an ELABORATION span cannot be deduced from truth-conditional semantics, consider the two 'decoupled' propositions ARRIVE*(princess)*∧WEAR*(princess, blue_chiffon_tracksuit.)* Applying our knowledge of linguistic conventions, the likeliest interpretation of this conjunction we would derive is an ELABORATION, where ARRIVE is the nucleus:

The princess arrived, wearing a blue chiffon tracksuit.

At first sight, it would seem this nuclearity decision follows directly from the propositions themselves. In event-contour terms, one proposition describes an achievement, while the other describes a state; the 'subjects' of the two are co-referential; achievements would tend to be considered more salient in narration than would states. These rationales describe how we might decide nuclearity here.[27] And yet, there would be nothing odd about a fashion columnist writing

The princess, who arrived next, wore a blue chiffon tracksuit,

which clearly reverses (and therefore defeats) the expected nuclearity.

As for the defeasibility of whether the span is an ELABORATION at all: YELLOW*(bird)*∧SHOOTS*(hunter, bird)* looks like an obvious attributive ELABORATION, satisfying M&T's criteria for the relation. As an ELABORATION, it would be textually realised as *the hunter shot the yellow bird*. But the spans could just as easily be related by CAUSE: *the hunter shot the bird* because *it was yellow (and he wanted to meet his colour quota.)*

2.1.6. Computational implications.

How is delimiting the purview of a rhetorical theory relevant to text generation in practice? Forbidding researchers like Rösner & Stede to come up with relations like UNTIL, designed to make their text generation more comprehensive, seems pedantic. After all, semantic distinctions like Temporality are as relevant as more orthodox rhetorical distinctions like MOTIVATION/ENABLEMENT, for a computer to decide what to say when. And these formal-semantic distinctions are as important as rhetorical theory itself in helping a computer establish which connective to use. Text generators use RST to decide these connectives at their planning stage; being told that a relation is merely a CIRCUMSTANCE will not help it decide between *until* and *during*.

True enough, none of my arguments need motivate changing the structure of text planners to distinguish between the formal-semantic and the rhetorical. Rather, setting such limits on RST ontology, as a theoretical issue, is important in making researchers *aware* that they are conflating different levels of analysis in text generation. If they are unaware of this conflation, they will end up constantly reinventing the wheel, and having unrealistic expectations of the explanatory and generative power of RST.

---

[27] In practice, of course, we would give a less formally stated account, like '*Arriving* is something you do, something that would matter in a story; no-one wants to focus on just wearing things.'

I believe this is what has happened with Rösner & Stede's UNTIL relation. A text generator does need to be able to distinguish between *until* and *during*. But that is not the business of a *rhetorical* theory; and phrasing the distinction as two more relations to add on to an unstructured list will not result in a properly comprehensive analysis. Temporal logic, not rhetorical theory, has been developed to analyse such distinctions; and researchers should exploit the results of temporal logic, not rhetorical theory, to understand and implement them. Implementing them *ad hoc*, purely in rhetorical terms, is self-defeating.

The relational criterion is also important in computational terms, in that it helps the generator decide where to expend disambiguation effort. If Volitionality is not a *relational* feature, but is local to the nucleus, then there is no use finding a distinctive connective to express Volitionality (I doubt any such can be found in English.) Instead, the generator should be focussing on the relation's nucleus, ensuring that the volitionality of the act it denotes is brought out by its wording. As Sanders *et al.* (1992) argue, this probably applies to the various subtypes of ELABORATION Hovy (1991) mentions as well.

2.2. What is rhetorically related?

A second problem, perhaps more immediately relevant to text generation, is the granularity of rhetorical analysis. What are the smallest units of text, between which rhetorical relations can be postulated? There are two extremes in the answer to this question. One extreme, subscribed to by text analysts, uses a syntactic definition: they consider the smallest unit to be the clause. The other extreme uses a semantic definition: they consider it to be the proposition. While it would be inaccurate to claim that either psycholinguists or computational linguists consistently maintain the latter position, both disciplines seem to lean more towards this mode of analysis than to the text-analytical mode.

2.2.1. Clauses.

Although M&T (1987) admit the size of the units their rhetorical analysis is built on is arbitrary, they assert that these units

> should be based on some theory-neutral classification. That is, for interesting results, the units should have independent functional integrity. In our analyses, units are essentially clauses, except that clausal subjects and complements and restrictive relative clauses are considered parts of their host clause units rather than separate units. (M&T:6)[28]

So if the text unit does not have 'independent functional integrity', it is not rhetorically interesting. Longacre (1983), for example, considers texts like *I know he's coming* as containing two rhetorically distinct text units, linked by an AWARENESS ATTRIBUTION. M&T conclude that

> The function of such attribution is in the domain of evidentiality with respect to the attributed material and thus is reasonably considered not as a distinct entity, but as part of the proposition that contains the attributional passage. (M&T:73)

In other words, the clause *he's coming* is not an independent clause that can be linked to *I know* rhetorically.[29] Rather, it is a complement of the clause *I know*, and as such is

---

[28] In their earlier paper, Mann & Thompson (1986) likewise state they are "primarily concerned with relations between clause-length units", although they admit the possibility that

> relations similar to the ones we will be discussing here can be found to hold between portions of text below the clause level, or between clause-level units and units below the clause level (Mann & Thompson 1986:59)

[29] Nor vice versa: *I know* isn't any more syntactically independent in this text. Even if the syntactic dependence is loosened, as in *He's coming; I know*, the semantic criteria M&T themselves use would exclude such a relation

(continued overleaf….)

subcategorised by *I know*. As a complement of this clause, it is bound more closely to its nucleus than a rhetorical theory would care to unravel.

That is a syntactic way of describing this dependency; in semantic terms, *he's coming* is an argument of another proposition. M&T do not find propositional arguments rhetorically interesting; for that reason, they exclude such complement clauses from their analyses. Indeed, if such clauses *were* included in a rhetorical theory, it would have to incorporate not only a formal semantics (as argued in the previous section), but also—since it would need to account for syntactic complementation—a case grammar.

By contrast, in the sentence *I'm rolling out the red carpet because he's coming*, the phrase *he's coming* is not a complement of *rolling out*, but an adjunct. As such, M&T would consider it as having 'independent functional integrity',[30] and they would analyse the adjunct–main clause relation rhetorically.

So in essence, Mann & Thompson's view of rhetorical structure is interclausal, although they give no *a priori* cause to discourage intraclausal analysis. They do give semantically-motivated reasons against analysing relations like QUOTATION and AWARENESS; as I have argued, these semantic reasons have clear syntactic correlates. To put it in more formal terms: Classical RST is *syntactically* delimited by the fact that rhetorical relations can hold between adjuncts and their heads at the S level, but not between complements or lower level adjuncts (like relative clauses), and their heads. The proscription involving adjunct level is not built in to the theory; the proscription against complements is.

In opposition to the classical view, there seems to be an overwhelming consensus amongst text generation researchers (de Souza *et al.* 1989, S&dS, Hovy 1990, Hovy 1991, Stede 1992, Rösner & Stede 1992, Vander Linden *et al.* 1992, Hovy 1993, Krifka-Dobeš & Novak 1993, Vander Linden 1993) that, in order to produce acceptable text, a text planner needs to have access to the *intraclausal* realisation of rhetorical relations. Since M&T have kept aloof from the intraclausal domain, several of these researchers go to great lengths to argue that this is necessary:

> […] Whereas Nuclei are only realisable as S-bars, Satellites may also be realised as sub-sentential structures (e.g. adjectives, noun-phrases or prepositional phrases.) In this respect we depart from the clause-combining proposal of [Matthiessen & Thompson (1988)], which suggests that satellites should also be realised as S-bars. Instead, we take the view that the semantic subordination of Satellite to nucleus should be expressible syntactically as embedding. (de Souza *et al.* 1989:228)

> Earlier work on RST has focused on the higher levels of text structure […] In our corpus, however, we found that most of the rhetorical complexity was located at the clausal and subclausal level. This required the extension of RST into the clause in order to provide a uniform treatment for clauses and subclausal relations such as prepositional phrases. (Vander Linden *et al.* 1992:184)

> [RST] assumes that text spans are of this minimal [clausal] size, because one cannot really apply RST-based planning to levels below that. Clearly the presence of constituents below the level of a clause, such as modifiers, should also play a role in planning text messages, and knowledge in a *real* knowledge base is certainly not always represented in clause-sized chunks. (Krifka-Dobeš & Novak 1993:103)

However, M&T had a reasonable linguistic constraint on what their analytic units were. What are the text units being invoked by the computational researchers? None of them gives an answer any more specific than 'propositions'. Presumably, none of them feel the

---

from RST (all the more so if, as Thompson herself has done recently (Thompson & Mulac 1991), we consider I know here merely a grammaticalising epistemic marker, and no longer an independent proposition.)

[30] Whether the verb is inflected or not makes no difference to whether a clause is 'functionally independent'; M&T analyse non-finite clauses like *while lubricating it at the same time* and *to see which Syncom diskette will replace the ones you're using now* as independent text units.

need to provide a *linguistic* characterisation of what their units are. These propositions are the basic units of their knowledge representation languages (typically either based on or isomorphic to first order logic.) The facts in their knowledge database have to be realised textually, and many of them can only be realised intraclausally.[31] The researchers probably don't feel the need to ponder too deeply any implications all this may have for RST itself as a theory.

In order to explain these implications, I now look more closely at what exactly the propositions in these systems are, and how they relate to text structure.

2.2.2. Propositions.

In a propositional analysis of text, anything that *can* be regarded as a proposition—that is, the application of a predicate to some arguments—*is*. This means that relative clauses and adjectives are regarded as distinct units. It also means that sentential complements, possessives, and nominalisations are all regarded as distinct units. To give an example: the following text would be analysed in Classical RST as two clauses linked in a CAUSE relation:

> *Since a sultan's harem might house 500 wives and concubines, only the best dancers could attract his attention.*

Fletcher (1981) proposes the following as a propositional analysis of the same text:

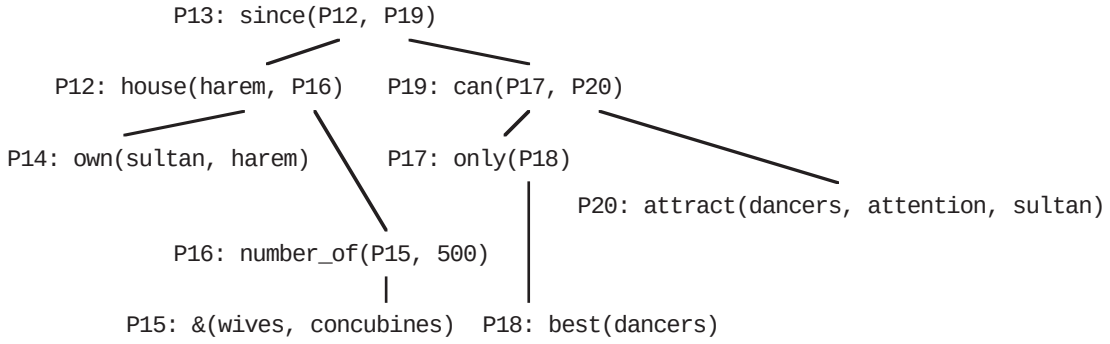

Fig. 2.1. Fletcher's (1981) propositional analysis of 'belly-dancing' text.[32]

This analysis incidentally illustrates that propositional analysis and clause-level analysis can still be congruent: the top level of the propositional analysis is a CAUSE relation between two propositions, which is exactly the same as the Classical RST analysis of the text. This points to a compatibility between syntactic, semantic, and rhetorical structure.

As another example of the depth of analysis possible in this paradigm, van Dijk & Kintsch (1983), in their discussion of the effect of proposition density on reading, identify no less than eight propositions in the following single clause:

> *Cleopatra's* ***downfall lay*** *in her* ***foolish trust*** *in the* ***fickle political*** *figures* ***of*** *the* ***Roman*** *world.*[33]

Various experiments have demonstrated the psychological reality of propositions. Graesser *et al.* (1980) have established that the number of propositions in a sentence is

---

[31] See Chapter 4 for further discussion.

[32] Indeed, this analysis underestimates the number of propositions in the text. The modals *could* and *might* are not accorded their own propositions, and pluralisation is not explicitly represented in the tree.

[33] I have set the words corresponding to the predicates of each of these propositions in bold type.

correlated with reading time, with an average of 117 ms required per proposition.[34] Thus, a sentence containing many propositions will take longer to read than a simpler sentence, even if they contain the same number of words.

So are these propositions the basic units of a rhetorical theory? They *are* the basic units of van Dijk & Kintsch's (1983) theory of how text coherence is built up in the mind. But their theory is not rhetorical: they make most use of other coherence devices, such as sharing discourse entities between propositions (e.g. anaphora.) And as seen, a theory which considers *I know he's coming* as two distinct rhetorical units (as a full-blown propositional theory must do) is not admissible according to M&T's perception of rhetorical theory.

It may still be possible to modify the syntactic distinction I made above, and admit all sentential *adjuncts* (at whatever syntactic level) into rhetorical theory, while excluding sentential *complements*. This would capture M&T's intuition that arguments of propositions cannot constitute independent rhetorical units, while drawing the obvious link between, say, non-restrictive relative clauses and ELABORATION sentences. (eg. *The prince, who was a complete doofus, bade the minstrels play* vs. *The prince bade the minstrels play. He was a complete doofus, you see.)*

On the other hand, such a theory might mandate making overly fine functional distinctions between text spans; for example, restrictive relative clauses would be excluded from the theory, while non-restrictive clauses would not. I will not pursue this line any further in this work, but it would be a very interesting avenue for research, to investigate how robust such a rhetorical theory might prove.

Rösner & Stede's (1992) analytic heuristics, the *paraphrase test* and the *gloss test*, could also prove useful for distinguishing between rhetorically relevant and irrelevant propositions:

> The 'paraphrase test' may help when a sentence is not immediately separable into 'basic units' for RST, e.g., when the status of a prepositional adjunct is debatable. Is there a reformulation for the same content that uses a subordinate clause or separate sentences? If so, there is strong evidence to treat the adjunct as the realisation of a proposition in its own right.
>
> The 'gloss test' is a kind of multilingual version of the 'paraphrase test'. When we analyzed English, German and—recently—French versions of manual paragraphs in parallel, we asked whether in cases of differing surface constructions 'glosses' of the other language's version would be possible as alternatives as well. If so, we were more convinced that the common RST analysis was well-founded. (Rösner & Stede 1992:200–201)

With their paraphrasing appraoch, Rösner & Stede are explicitly appealing to the functional equivalence of text spans: if a paraphrase with a different rhetorical realisation conveys the same function, then the rhetorical realisations are equivalent at some underlying level. The problem is that paraphrasing is concomitant with some semantic change; judging two paraphrases to be functionally equivalent is rather subjective, and difficult to reproduce.

The distinctions made in these approaches, between rhetorically relevant and irrelevant propositions, would give us a linguistically well-motivated theory of text structure, which researchers in text generation can work with. However, such a theory would have several shortcomings Classical RST has managed to avoid, as I discuss below.

---

[34] In contrast to earlier research, Graesser *et al.* isolated the effect of number of propositions from that of other factors correlated with it—such as introduction of new referents, and how 'narrative' the text is.

2.2.3. Problems with a propositional RST.

Rösner & Stede (1992) report their experiences in using RST in the computer generation of manuals in both English and German. They make a number of suggestions to alter the structure of Classical RST, in order to bring it in line with the requirements of text generation. They also point out a number of problems they encountered with RST.

The problem I wish to concentrate on arises from their analysis of the following text:

> *[The spark plugs must be securely tightened]$_{8a}$, but [not over-tightened]$_{8b}$. [A plug that's too loose]$_{9a}$ [can get very hot]$_{9b}$ and [possibly damage the engine]$_{9b}$; [one that's too tight]$_{10a}$ [could damage the threads in the cylinder head]$_{10b}$.*[35]

They give the Classical RST analysis of this text as shown in Fig. 2.2.

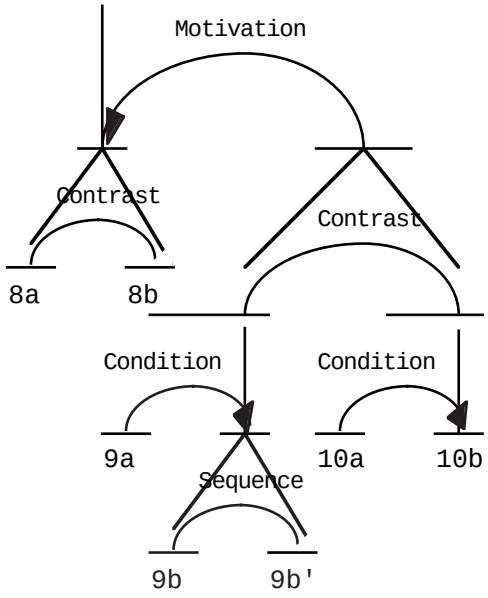

Fig. 2.2. Rösner & Stede's version of a Classical RST analysis of the 'spark plugs' text.

Rösner & Stede argue that the adjacency principle of RST (which requires that only adjacent text spans should be rhetorically linked) renders the analysis in Fig. 2.2. unsatisfactory. In particular, they claim spans 8a–9 and 8b–10 are what are truly being contrasted, with 8a acting as the MOTIVATION for 9, and 8b for 10. They give the following rephrasing of the text to convey these dependencies:

> *[The spark plugs must be securely tightened]$_{8a}$, because [a plug that's too loose]$_{9a}$ [can get very hot]$_{9b}$ and [possibly damage the engine]$_{9b}$. [On the other hand it should not be over-tightened]$_{8b}$, because [a plug that's too tight]$_{10a}$ [could damage the threads in the cylinder head]$_{10b}$.*

They conclude that, since this analysis violates the adjacency principle for the original text, RST is inadequate to represent the full complexity of rhetorical relations in a text, which actually constitute

> not a simple tree but a net of related propositions[;] and some of these relations are no longer represented in the RST tree (Rösner & Stede 1992:203)

---

[35]The bracketing of the text is Rösner & Stede's, and is somewhat loosely used: they bracket *a plug that's too loose*, a noun phrase, as a rhetorical unit; presumably, they intended to separate the relative *clause, is too loose*, from the rest of the clause.

Clearly, though, such a violation of the adjacency principle only arises because Rösner & Stede are using a version of RST with much finer units than Classical RST tends to use. An M&T-style analysis would be quite contented with the level of analysis in Fig. 2.2. If anything, it would probably eliminate the distinction between 9a and 9b, and 10a and 10b, as distinct rhetorical units—giving the analysis in Fig. 2.3.

The reason a classical analysis would not be perturbed by Rösner & Stede's observations is that the clause, and at times even the sentence, are treated as atomic units in a classical analysis—which sees no reason to delve into these units any further. There is no question that there is a cohesive link between spans 8a and 9a. It it also obvious that, if we considered *all* the cohesive links in a text, rather than just those links RST concentrates on, we would end up with the kind of tangled net Rösner & Stede speak of, rather than neat trees.

But then, that is precisely why Classical RST disregards all but a small number of cohesive links. An intersentential analysis of coherence, treating sentences as impenetrable atoms, is likeliest to give a neat tree analysis, rather than a more complex cohesive net. An interclausal analysis, such as Rösner & Stede have presented, is less likely to do so; and a functional-propositional analysis, advocated by researchers like Scott & de Souza in their treatment of ELABORATION,[36] is least likely of all.

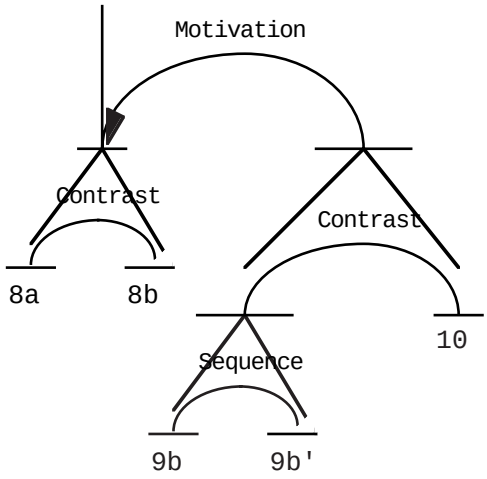

Fig. 2.3. An M&T-style Classical RST analysis of the 'spark plugs' text.

The 'neat tree' approach of Classical RST gives results, is effective in accounting for text structure in a broad sense, and doesn't allow the analyst to 'miss the forest for the trees'. In text-generation terms, it can be much more straightforwardly applied to the text planning task than a messier, cohesive approach, which would probably do much superfluous work (since a lot of text cohesion will arise as a by-product of classically-oriented rhetorical text planning anyway.)

On the negative side, Classical RST does not give a full account of coherence. But its advantages outweigh its disadvantages significantly enough in its primary fields of application (creating cohesion in computer-generated texts, and analysing the coherence of essentially clause– and sentence–sized units in text analysis), that it is not worth jettisoning.

[36] See Chapter 4 for a fuller discussion.

2.3. Conclusion.

As an analytic theory, the use of RST is straightforward: analysts have an intuitive sense of what linguistic elements are being related, and of how semantically detailed those relations are allowed to be. When in doubt, they can appeal to the worked examples in M&T. Rhetorical analysis seems to proceed on an 'everybody knows that' basis.

But the transposition of RST into a computational domain, and the rather different demands made on it by text generation researchers, has highlighted the fact that some crucial ontological parameters of RST are underspecified, or inappropriate for computational applications. This is not as much a problem for the computationalists (who have empirically worked their way around these problems) as for the theoretical linguists, now confronted with the fact that some aspects of RST have not been formalised sufficiently for the theory to be formally sound and reproducible.

I have addressed two particular problems: what types of relations should a rhetorical theory distinguish between, and what types of linguistic elements should a rhetorical theory relate.

The first issue decides the linguistic orientation of the theory. If the theory makes the types of distinctions a formal semantics would make, the theory becomes essentially formal-semantic; if it makes the types of distinctions a pragmatic theory (particularly Speech Act Theory) would make, then it becomes a pragmatic theory.

This issue also decides how analytically useful the rhetorical theory is, and in what domains. If a rhetorical theory makes all the discriminations a formal semantics makes, then in many respects it becomes indistinguishable from formal semantics—in particular, it may have access to lexical-level meaning, making it a very different kind of theory from what we might expect. If, on the other hand, it becomes pragmatics-driven, it will tend to move away from both the clausal level of analysis (towards a more illocutionary level), and formal-semantic–motivated distinctions.

As an attempt to answer this challenge, I consider Sanders *et al.*'s (1992) relational criterion, used to determine whether a particular meaning feature should be used to discriminate between relations in a rhetorical theory. The criterion gives an essentially pragmatic answer to this question, although it admits semantic distinctions to a sufficient extent that the descriptive adequacy of M&T's rhetorical relation inventory is not seriously challenged. The criterion is not without problems; it seems to reject too many features as worded by Sanders *et al.* (although I have modified it to account for, for example, Hallidayan textual meaning (discourse motivations)); and its view of formal-deductive semantics is somewhat simplistic. But contrasted to the complete open-endedness that has prevailed in RST to date, the criterion is an important advance in making RST a more rigorous theory.

The second issue—what types of linguistic elements should a rhetorical theory relate—creates a split between discourse linguists, computational linguists, and psycholinguists. Text linguists take a syntactic view of what the atoms of rhetorical theory should be, using clauses. Psycholinguists are not involved in rhetorical theory as distinct from coherence theory in general; but in coherence theory, they take an explicitly semantic approach to the issue, using propositions. Motivated both by psycholinguistic work and by their own knowledge representation schemes, computational linguists lean towards a propositional approach; at the very least, they analyse syntactic units much smaller than discourse linguists would analyse (such as relative clauses and adjectives.)

In my analysis, I attempt a more precise syntactic description of both the discourse-oriented approach, and of what an approach more consistent with computational demands might look like. I also briefly consider why a rhetorical theory with finer syntactic units

than is traditional (such as my formalisation of the computational approach to RST) might encounter analytical problems the more traditional analysis has managed to avoid.

My proposal for formalising the syntactic ontology of computational RST does not purport to settle the issue in the way the relational criterion does; it is much more tentatively stated. But I am hopeful that research like this will encourage text generation workers to look more closely at their assumptions on what constitute the units of their RST analyses—and to approach their analyses with more linguistic rigour than has been the case. Computational approaches which integrate linguistic analyses more closely (rather than piecemeal and *ad hoc*) are sure to produce linguistically sounder output.

3. RST Taxonomy.

As argued in Chapter 1, there is a problem with researchers—particularly computational linguists—taking the inventory of rhetorical relations presented in M&T as an unstructured block. As a result, text planners have to choose the rhetorical relations they will exploit out of an unstructured list, which is intrinsically difficult to formalise. As Hovy admits, functionalist theories of text structure like RST

> tend to be weakest on the formal aspects of the discourse structure and the formal definitions of the operators (for example, no adequate language has yet been developed for defining the segmentation operators' contents, nor has a well-specified formalism been developed for capturing discourse structure.) (Hovy 1993:59)

By using such an unstructured list of relations, researchers also fail to account systematically for the range of meaning of ambiguous connectives: resemblance in meaning between relations may be reflected by shared connectives. For example, the connective *because* is shared between MOTIVATION, PURPOSE, JUSTIFY, EVIDENCE, and CAUSE.

Attempts to taxonomise rhetorical relation inventories are frequently played down by researchers working on text coherence, both within the RST tradition and outside it:

> Several people have suggested that we create a taxonomy of the relations in order to present the important differences among them. However, no single taxonomy seems suitable. Depending on one's interests, any of several features could be made the basis for grouping them. (M&T:17)

> There is no single, uniquely correct inventory of the types of conjunctive relation; different classifications are possible, each of which would highlight different aspects of the facts. (Halliday & Hasan 1976:238)

Indeed, M&T have *deliberately* made RST taxonomically open-ended:

> Despite our efforts to say the opposite, some have read our other papers as suggesting that the relations are a closed list, a kind of one-dimensional text theory. We see it as an open set, susceptible to extension and modification for the purposes of particular genres and cultural styles. (M&T:48)

As I argued in Section 1.5., this kind of open-endedness may buy some descriptive flexibility, but at the cost of reproducibility of analyses and theoretical coherence.

Notwithstanding such opposition, in the work that *has* been done to taxonomise rhetorical relations, there seems to be a recurring pattern of classificatory parameters. I believe these parameters can be integrated into a workable whole. To this end, I will evaluate what I consider the major such research efforts (Longacre (1983), Halliday & Hasan (1976), Hobbs (1985), Mann & Thompson (1986, 1987), Knott & Dale (1992), Sanders *et al.* (1992), and Maier & Hovy (1993)); my yardstick for comparison will remain M&T's 1987 inventory. While much of this work is outside the RST framework proper,[37] it analyses the same text-linguistic phenomena as RST; and the resulting analyses are congruent enough to RST (hierarchical text structures, same types of rhetorical relations ascribed, same basic units of text analysis) for comparison to be valid and instructive.

---

[37]Longacre (1983), Halliday & Hasan (1976), Hobbs (1985), Knott & Dale (1992) and Sanders *et al.* (1992) are outside the RST tradition, although the last research effort acknowledges the descriptive adequacy of RST. Maier & Hovy (1993) is more closely derivative of RST, Mann & Thompson (1986) is essentially an early version of RST, and M&T (1987), of course, define RST itself.

To be useful for the requirements outlined in Chapter 1, I believe a rhetorical taxonomy should satisfy three criteria. It should be comprehensive, covering the entire rhetorical inventory. It should be externally motivated by some independent classificatory criteria, rather than empirically derived; otherwise, it is not falsifiable, and thus not scientific.[38] And it should be feature-based, to allow classes of relations and generalisations about relations to be formulated easily and flexibly, along more than one taxonomical dimension.

None of the taxonomies I consider satisfies all these criteria. Therefore, I propose my own synthesis of these taxonomies, which I feel does satisfy them all. In particular, the distinctions made between relations in my taxonomy are extensively argued for in terms of an external theory (advanced formal logic.) As a test of the explanatory power of this taxonomy, I apply it to several rhetorically ambiguous connectives, in order to account for their ambiguity systematically.

3.1. Longacre's taxonomy.

Longacre first wrote on the subject of rhetorical relations in 1971; his work thus predates most others' work in the field. Mann & Thompson admit that his work was influential in their formulation of RST. The 1983 version of Longacre's taxonomy is briefly outlined in Fig. 3.1.; it is more fully described in Appendix B.1.

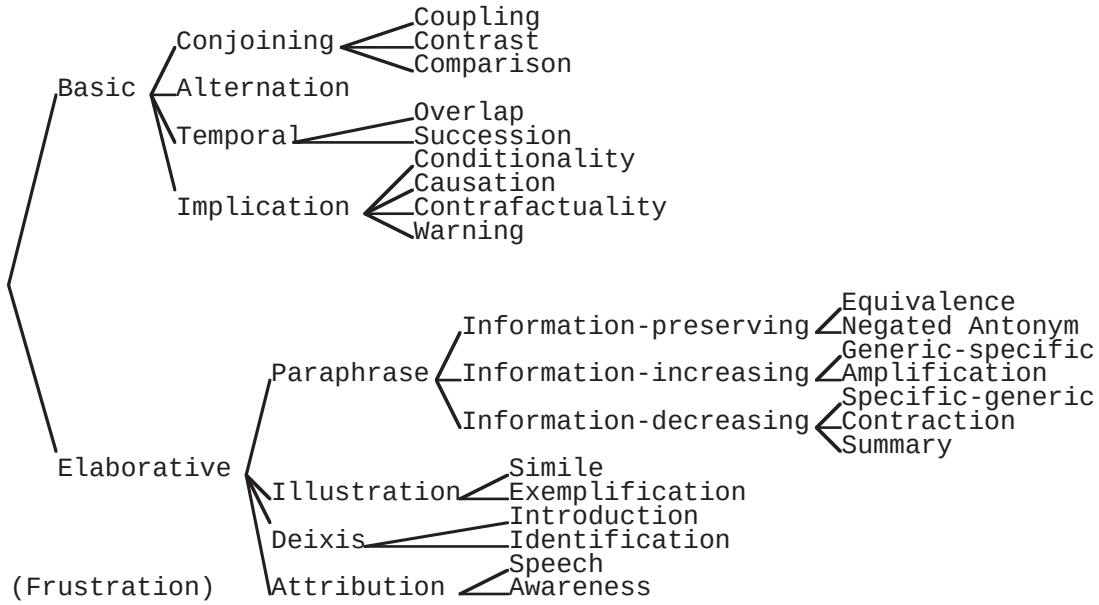

Fig. 3.1. Longacre's rhetorical taxonomy.

The Longacre taxonomy is an excellent and thorough description; unusually for such work, it is also cross-linguistically informed. Longacre is particularly astute in making a distinction between *Basic* and *Elaborative* relations.[39] Some of the taxa without RST equivalents are problematic because Longacre differs from M&T in what he considers the scope of a rhetorical structure theory (see Chapter 2.) For example, he includes quotation and evidentiality within the scope of his theory. M&T counter that

---

[38]It is true, however, that any proposed taxonomy, however deductivist or rationalist, still has to be subject to empirical testing on actual texts.

[39]This distinction has been subsequently used by Halliday, and is thence invoked in Matthiessen & Thompson (1988:298)

> [t]he function of such attribution is in the domain of evidentiality with respect to the attributed material and thus is reasonably considered not as a distinct entity, but as part of the proposition that contains the attributional passage. (M&T:73)

Similarly, M&T would not classify comparisons like *Gilgamesh is taller than Rostâm [is]* as rhetorical relations; they would class the 'clauses' here as complements of a proposition, rather than as rhetorically independent units.

M&T draw attention to differences between their work and that of Longacre and his colleagues. Despite these differences, Longacre's classificatory scheme can still be superimposed onto the RST inventory without much stretching of the framework.

As an example of these differences, M&T point out that the work of their predecessors considered rhetorical structure as linear, positing relations only between individual sentences and "successive combinations of clauses". Since RST is hierarchical rather than linear, they argue, RST includes relations that only appear between clause groups, rather than clauses. By binding just single clauses to their rhetorical structures, Longacre and his colleagues would ignore such relations.

But this distinction doesn't seem valid. SOLUTIONHOOD, EVALUATION, and the other relations missing from Longacre (1983) *can* be exemplified by a link between two single clauses (e.g. *I'm hungry. Let's go to Fuji Gardens.)* True, more prototypical instances of, for example, SUMMARY and ENABLEMENT obtain at higher levels of an RST tree; and BACKGROUND doesn't sit comfortably in Longacre's taxonomy. This does not make the two theories incompatible, though—particularly since Longacre taxonomises *Elaboration* relations (particularly his *Paraphrases)* in a discourse-based way compatible with RST—in contrast to his *Basic* relations, where logical connectives are used.

This brings up M&T's second objection, that Longacre taxonomises "using semantic principles of clause combining rather than in terms of discourse-organising principles," whereas RST gives "a functional model that asks what the text is doing for the writer." (M&T:40) At first blush, Longacre's usage of formal semantics—particularly of logical connectives—seems to contradict the view of rhetorical theory held by Mann & Thompson, and enshrined in the relational criterion. If the distinctions between Longacre's *Basic* relations are expressed in propositional logic, what place do they have in rhetorical theory?

But this would be to misunderstand the relational criterion. The criterion says nothing about whether a relation can be *expressed* in formal logic or not.[40] Rather, it asks whether the presence or absence of a particular rhetorical feature can be *deduced* directly from formal semantic information. CONDITION, for example, can be expressed quite easily in formal logic terms by the logical connective *if*.

Understandably, in a perlocution-oriented theory like RST, M&T do not wish to invoke formal logic in their definitions. But this does not make the formal logic paraphrase of CONDITION any less valid; and even the most fanatically pragmaticist rhetorical analyst would not exclude CONDITION from the rhetorical inventory, simply because it can be expressed in formal logic terms. Indeed, there is no real difference between M&T's definition of CONDITION and Longacre's more explicitly logical definition—as the following definitions show:

> ***(a)*** **Constraints on S:** S presents a hypothetical, future, or otherwise unrealized situation (relative to the situational context of S.)

---

[40]This is why, in Section 3.3., I do in fact use formal logic to express the effects of RST relations.

> **Constraints on the N+S combination:** Realization of the situation presented in N depends on the realization of that presented in S.
> **Effect:** R recognizes how the realization of the situation presented in N depends on the realization of that presented in S. (M&T:65)[41]

> ***(b)*** Hypotheticality includes certain unweighted *if* notions. This notional structure expresses a condition which implies nothing as to factuality of either member of the condition. It simply states a relation between an antecedent and a consequent, i.e., the consequent does not follow unless the condition stated in the antecedent also holds. (Longacre 1983:101)

In fact, there is nothing in Longacre's subdivisions of *Basic* relations (Conjoining, Alternation, Temporal, and Implication) which is alien to RST (though M&T eventually rejected DISJUNCTION (Alternation) as a distinct rhetorical relation.)

M&T's last objection, that a "theory of writers' goals is necessary for understanding the rhetorical organization of a text," is the most telling. Longacre's scheme is innocent of the *Presentational/Informational* distinction. As a result, some Presentational RST relations do not figure in it at all, and those that do (EVIDENCE, CONCESSION, BACKGROUND, ANTITHESIS) do not fit in their niches comfortably—they show up more by accident than by design.

To conclude: Longacre's is a coherent and well-ordered taxonomy, which makes several astute distinctions between relations, particularly in *Paraphrases*. It suffers from two major problems, however. It has an overly broad view of rhetorical ontology—which makes it include relations like SPEECH ATTRIBUTION, that most rhetorical analysts would not consider rhetorical; and by ignoring the functionalist approach to analysing text, it poorly covers more pragmatically-oriented relations.

3.2. Halliday & Hasan's taxonomy.

In their influential work, Halliday & Hasan (1976) outline a taxonomy of what they term 'Conjunction'. *Conjunction* is simply the form of cohesion which "resides in an abstract relation between one proposition and another." (Halliday & Hasan 1976:13)

The nature of *cohesion* itself is somewhat elusive in their work. They ostensibly define cohesion as a *semantic* property of a text, which obtains when the interpretation of some element in the discourse presupposes another. Cohesion, in turn, is realised by lexical ties which make the link between the two elements explicit.

What Halliday & Hasan do not make explicit is whether cohesion can be posited when no lexical tie is present. They imply it can be; yet their analysis seems oriented towards a classification of conjunctions, and ignores relations like SOLUTIONHOOD which are intrinsically hard to signal textually. In other words, they don't make the distinction between *coherence*, a semantic property, and *cohesion*, a textual property.[42] By considering connectives inessential to establishing rhetorical relations, RST is a coherence theory. Despite their metalanguage, Halliday & Hasan's is not: it is restricted to cohesion.

---

[41] A reminder: R stands for the reader (or listener) of the text; W for the writer (or speaker); N for the nucleus of a rhetorical relation; and S for the satellite.

[42] As Hoey's (1991:11) discussion shows, Halliday & Hasan and their successors tend to blur the distinction between the two. The earliest work Hoey mentions as making explicit the fact that coherence is possible in the absence of any cohesive link (and therefore, that coherence and cohesion are not synonymous) is Widdowson (1978), although Hasan (1984) explicitly defines coherence as a non-absolute property. The most useful definition distinguishing the two is by Beaugrande & Dressler (quoted in Hoey (1991)):

> "[Cohesion] concerns the ways in which the components of the surface text, i.e. the actual words we hear or see, are *mutually connected within a sequence*, [coherence] concerns the ways in which the components of the textual world, i.e. the configuration of concepts and relations which *underlie* the surface text, are *mutually accessible and relevant*." (de Beaugrande & Dressler 1981:3–4)

Another complication is that Halliday & Hasan's atomic unit of analysis is the sentence. The authors do not consider subordinate clause linking in their work at all. For example, though their taxonomy is pretty much a taxonomy of connectives, rather than a taxonomy of relations, they do not even consider *if* a marker of CONDITION (since *if* does not operate intersententially.) Instead, they name only the intersentential connectives *then*, *in that case*, *in such an event*, and *that being so* as markers of this relation.

Halliday & Hasan's scheme is briefly outlined in Fig. 3.2., and described more fully in Appendix B.2.:

	External/Internal	**External**	**Internal**
Additive	Simple Additive Negative Alternative		Complex Apposition Expository Exemplificatory Comparison Similar Dissimilar
Adversative	Adversative 'proper'	Contrastive	Contrastive Correction Dismissal
Causal	General Specific Reason Result Purpose Conditional		Reversed Specific Reason Result Purpose Respective
Temporal	Conclusive	Simple Sequential Simultaneous Preceding Complex Immediate Interrupted Repetitive Specific Durative Terminal Punctiliniar	Internal Temporal 'Here and now' Past Present Future Summary Summarizing Resumptive)

Fig. 3.2. Halliday & Hasan's conjunctive cohesion taxonomy.

Halliday & Hasan's scheme is something of a grab-bag. It isn't as well ordered as Longacre's scheme—their categories are much more open-ended. It has minimal coverage of Longacre's *Elaboration* relations—a natural consequence of the fact that these relations seem intrinsically hard to signal textually.[43]

In addition, the distinctions made between connectives in their scheme are not always informative about their underlying rhetorical relations. For example, the split of *Cause* into *General* and *Specific* is misleading. There are not two different types of causation at work here. Rather, the Specific connectives *for this reason*, *as a result* and *for this purpose* specify which of the three possible causations holds: CAUSE, RESULT or PURPOSE. The connective *so*, on the other hand, is ambiguous between the three, and so is given a separate, *General* class. The distinction between 'emphatic' and 'simple' markers is also not rhetorically informative, since the underlying rhetorical relation signalled remains the same. So Halliday & Hasan are taxonomising *connectives*, and not cohesive *relations*.

[43]See Chapter 4 for a fuller discussion.

The most important contribution Halliday & Hasan make is introducing the *Internal/External* distinction. This distinction is reminiscent of M&T's *Presentational/Informational* distinction, but isn't quite the same. As M&T point out, RST is a relational theory. Nuclear and satellite text spans are never considered in isolation (unlike e.g. McKeown's Schemas), but always in terms of how they *combine* to produce an effect on the reader. Thus, the *Presentational/Informational* distinction describes whether the relation *as a whole* is intended to inform the reader, or to increase some inclination in them.

The metalanguage of Halliday & Hasan, on the other hand, as with that of Longacre, is bound up with logical connection. The *Internal/External* distinction considers what elements are bound by the connective—whether they are locutions or illocutions.[44] So the Internal/External distinction scrutinises the pragmatic status of the nucleus and the satellite separately, rather than describing their effect in combination.

Thus, PRESENTATIONAL SEQUENCE[45] is not associated with any temporal sequence of events denoted by the utterance—unlike a proper SEQUENCE relation. Rather, the sequence is one of illocutionary acts. Similarly, EVIDENCE must be considered, not as a causal relation between events (like CAUSE), but as a causal relation between mentioning an event or state (the *Argument)* and an illocutionary act (the *Claim.)*[46]

The appeal of this distinction is that it captures the differences between such relations as EVIDENCE and CAUSE, or PRESENTATIONAL SEQUENCE and EVENT SEQUENCE, while allowing us to retain the intuitive links between those pairs (that they are both Causal and Temporal, respectively.) But there are two problems with using it as a classificatory parameter.

The first is that the distinction between Internal and External relations becomes less clear-cut as one moves across taxonomical parameters. Halliday & Hasan report that the distinction is "fairly easy" to make within Temporal relations, but much more elusive when it comes to Causal, Adversative or Additive relations. Sanders *et al.* (1992) discover a similar phenomenon with their Negative relations (corresponding loosely to Halliday & Hasan's Adversatives); I will discuss this more fully in section 3.4.

The Internal/External distinction seems hard for people to unravel while analysing text (as indicated by Sanders *et al.*'s experimental data); and rhetorical inventories like M&T's ignore the distinction entirely in places. So is there any point in making this distinction a uniformly applicable parameter, creating seemingly artificial rhetorical distinctions?

Perhaps there is, if RST is to help generate formal semantic models of discourse; it may even turn out to be of use in computational linguistics. But the distinction doesn't seem essential in all cases to explaining human linguistic competence—which is the function of RST as a text-linguistic framework. In other words, if people don't think a Internal/External distinction is anything more than hairsplitting when it comes to Adversative rela-

---

[44]The distinction wasn't expressed in those explicit Speech Act terms by Halliday & Hasan, nor by any of their predecessors they cite. But the terms they *do* use leave little doubt as to what is intended, and these terms are used explicitly by Sanders *et al.* (1992:6). One of the earliest discussions of this distinction, Rutherford (1970), firmly places it in the context of the then current Performative Hypothesis, which attempted to account for the distinction between locutions and illocutions in deep-structure syntactic terms.

[45]A relation M&T considered for inclusion in their inventory, but did not put in their final draft. Eg. ***First**, you taped a counter-revolutionary song. **Second**, you had announced you were doing so, against the advice of your elders*.

[46]But see Sweetser's (1991) account of EVIDENCE, described in section 3.7.1.

tions, this is probably for a good psycholinguistic reason, and we needn't insist on the distinction in that part of the taxonomy.

The second problem with the Internal/External distinction is a confusion between it and M&T's Presentational/Informational distinction. An Internal relation needn't be Informational, nor vice versa. The Internal/ External distinction describes the elements entering into a connective rhetorical relation; the Presentational/Informational distinction describes the functional effect of the entire rhetorical predication. The two are not always correlated.

CONCESSION, for example, is considered by M&T a Presentational relation; yet Sanders *et al.* (1992) posit an Internal and an External relation (to use their terminology, Semantic and Pragmatic), both of which correspond to CONCESSION. Since we are attempting to list points in common between relations, both classifications can be exploited, even if they conflict with each other. But if a *general* RST taxonomy is to be posited, which is more appropriate? I will return to this question in section 3.7.

3.2.1. Knott & Dale's substitution classes.

Knott & Dale (1993) are motivated by a feeling that RST requires a taxonomy (appealing to arguments similar to those I have presented.) As a result, they approach the question of building up a taxonomy by looking directly at the linguistic evidence provided by textual connectives. Their analysis thus lies in the same domain as Halliday & Hasan's (1976): it does not account for relations that have no corresponding textual marker, and the distinctions made in their scheme are only those actually reflected in connectives. However, they circumvent some of the shortcomings of Halliday & Hasan's scheme by imposing a more soundly motivated, hierarchical taxonomy on their data. They do this by using substitution classes of connectives.

To illustrate these substitution classes: the connectives *in the end*, *eventually* and *ultimately* are all considered intersubstitutable in a text. That is, they can substitute for each other in a discourse, without altering the sense of the text, or the relation signalled by the connective. Knott & Dale call this substitution class LAST STEP IN TEMPORAL SEQUENCE. The expressions *above all* and *most importantly* form another such substitution class: LAST STEP IN PRESENTATIONAL SEQUENCE. Furthermore, *finally* can replace members of both substitution classes, without changing the meaning of the sentence; the reverse is not necessarily true.

Thus, the meaning of the sentence *Eventually he went* can also be expressed by the sentence *Finally he went*; but *Finally he went* does not mean the same as *Above all, he went*. Therefore, *finally* belongs to its own substitution class, LAST STEP IN SEQUENCE, of which the two other substitution classes are considered children (more specific instances), in a hierarchical taxonomy of connectives.

As a result of their analysis, Knott & Dale identify six classes of connectives, such that no connective in any one class is intersubstitutable with a connective in any other. The classes are:

- Sequences (cf. SEQUENCE) e.g. *to start with, next, in the end*;
- Event Situations—temporal and spatial (cf. CIRCUMSTANCE) e.g. *after, until, wherever*;
- Causal (cf. CAUSE, CONDITION, OTHERWISE) e.g. *because, so, if, otherwise*;
- Similarity (cf. JOINT) e.g. *also, once again, as well*;
- 'Negative Polarity' (corresponding to Longacre's Adversatives; cf. CONTRAST, CONCESSION, ANTITHESIS) e.g. *but, nevertheless, although*;
- Clarifications.

The last class, *Clarifications*, contains several subclasses which are not intersubstitutable—in contrast to the other classes. The subclasses include RESTATEMENT *(that is)*,

POSITIVE RESTATEMENT OF A NEGATIVE *(on the contrary),* EXEMPLIFICATION *(for example),* ADDITIONAL INFORMATION *(in fact),* and SUMMARY *(in short.)*

As suggested by the RST relations each class is compared with, Knott & Dale's connective classes are compatible with rhetorical relation schemes like RST. However, Knott & Dale do not aim to set up their own rhetorical scheme, but to provide a pre-theoretical framework for linguistically motivating such schemes. The major importance of their work is therefore as a test-bed for such schemes; and I use it in the remainder of this chapter to support my argumentation on rhetorical categories.

3.3. Mann & Thompson's classifications.

Mann & Thompson are sceptical about any attempts to impose an overall taxonomy on RST. But their work is full of one-off usages of parameters, to establish some order in the inventory—particularly when the distinction between certain pairs of relations becomes too close to unravel unaided. The parameter that has received the most attention so far is the Presentational/Informational distinction. But M&T have recognised structure elsewhere in their inventory as well. Here follow some instances of such structure.

3.3.1. The CAUSE Cluster.

M&T (1987:57) consider VOLITIONAL CAUSE, NON-VOLITIONAL CAUSE, VOLITIONAL RESULT, NON-VOLITIONAL RESULT, and PURPOSE—relations which all involve 'notions of cause'—as a cluster. They subdivide this cluster according to the binary features of Nuclearity (a syntactic factor[47]), and Volitionality (a semantic factor); PURPOSE is considered "definitionally neutral" across these parameters. (That is, it can be either volitional or non-volitional, and the nucleus can be either the logical antecedent or the consequent.)

Being grouped together so overtly, these relations seem to be semantically closer to each other than are any other RST relations. As I argued in Chapter 2, both Nuclearity and Volitionality are parameters which should be regarded as irrelevant in a rhetorical taxonomy. So in taxonomical terms, the CAUSE and RESULT relations in the cluster can be collapsed to the single relation CAUSE. This is in fact how I treat these relations throughout this thesis.

What remains of the CAUSE cluster, in that case, is the notion that CAUSE and PURPOSE are closely related. I see no motivation, however, for also reducing PURPOSE to CAUSE. Unlike the Volitionality distinction, the two relations have distinct linguistic realisations *(because* vs. *to* clauses in English.) Furthermore, the distinction passes the relational criterion: Using only referential semantic information, it is impossible to tell whether the propositions RUSH_INTO*(alex, bathroom)* and MISS*(alex, show)* are related as a PURPOSE span *(Alex rushed into the bathroom so he could miss the show [because no-one would drag him out of there])* or as a CAUSE span *(Alex rushed into the bathroom because he was going to miss the show [so he wouldn't be late by taking too long].)*[48]

---

[47] As I have argued in sections 1.4. and 2.1.2., a shift in nuclearity is not a relational difference *per se*, since it does not give rise to a distinct rhetorical predicate, but merely to the reversal of arguments in a rhetorical predication. Apart from the resulting shift in discourse-focus, CAUSE and RESULT are not considered as having any essential difference for the purposes of this thesis.

[48] Normally, the temporal properties of the satellite are a good clue as to which of the two relations holds. A satellite that temporally precedes its nucleus would typically be a CAUSE; e.g. *Alex rushed into the bathroom because he had left his toothbrush there*. A satellite that temporally follows its nucleus, or at least is unrealised at the time of the nucleus, would typically be a PURPOSE; e.g. *Alex rushed into the bathroom to get his toothbrush*.

Whether temporal distinctions matter to the relational criterion is not settled, as I discussed in Chapter 2. But the examples given in the main text shows this point is moot: both satellites there are in the future of their

(continued overleaf....)

3.3.2. The Causation cluster.

Mann & Thompson (1986) concede that EVIDENCE, JUSTIFY, MOTIVATION and REASON[49] "have proven easy to confuse". To aid discrimination between them, they create a mini-taxonomy of these relations. The intuitive link between all these relations is a notion of causality, although for three of them the causation is clearly at an illocutionary level (whereas for the CAUSE cluster proper, it was confined to the real-world level.) Since these relations can be confused, they form a cluster; for convenience I will call this the Causation cluster to distinguish it from the CAUSE cluster, which is organised around CAUSE.

The distinctions Mann & Thompson posit between the relations are as follows:

> EVIDENCE is discriminated from the others on the basis that all of the others address an action, but EVIDENCE does not. JUSTIFICATION is discriminated from the remainder in that it addresses a speech action found within the same text, while the others do not. MOTIVATION addresses a potential action, whereas REASON does not. (Mann & Thompson 1986:62)

The distinctions are phrased in vague terms (like *address.)* The distinctions also seem to violate the relational criterion, being constrained to differences in the event status of their respective satellites, rather than ranging across the entire relation. But a look at the effects M&T ascribe to each of these relations (as explicit definitions) should help clear up the confusion:

Relation	**Effect**
EVIDENCE	R's belief of N is increased.
JUSTIFY	R's readiness to accept W's right to present N is increased.
MOTIVATION	R's desire to perform action presented in N is increased.
VOLITIONAL CAUSE	R recognizes the situation presented in S as a cause for the volitional action presented in N

Table 3.1. Effects of relations in the Causation cluster.

It is possible to recast these 'effects' into symbolic logic as causal consequents (states intended to be caused by W's locution)—even though this may be against the spirit of M&T, who objected to Longacre's use of formal logic. Using symbolic logic should make the relationships between these rhetorical relations more explicit.[50] Applying symbolic logic to the effects in Table 3.1., we can derive Table 3.2.:

---

nuclei, so a referential semantics would be unable to distinguish between them. The actual distinction between CAUSE and PURPOSE—intentionality—is not manifest in the individual text spans, and presumably would not be included in the type of referential semantics Sanders *et al.* (1992) had in mind when proposing the relational criterion.

[49]This relation seems to correspond to the VOLITIONAL CAUSE M&T subsequently defined in 1987, since non-volitional examples in Mann & Thompson (1986) are listed separately under CAUSE

[50]As mentioned above, using formal logic to *describe* rhetorical relations, rather than deduce their presence, does not violate the letter of the relational criterion. I do not believe it violates its spirit either. The problem the relational criterion is intended to resolve is rhetorical theory becoming as strong as truth-conditional semantics (which would unnecessarily complicate rhetorical analysis) or as weak as truth-conditional semantics (which would render it powerless to cope with the illocutionary and perlocutionary aspects of RST).
The symbolic logic I use is not in fact truth-conditional, since it incorporates deontic and belief logic: it is capable of expressing illocutions such as directives, and perlocutions such as belief, which propositional logic by itself cannot express. In any case, using formal semantics to attempt to model rhetorical relations is not a problem for rhetorical theory, but for formal semantics—since it is formal semantics I am attempting to fit to the task, not vice versa.

Relation	**Logical Paraphrase of causal consequent**	**Translation of Logical Paraphrase** *(cf. Effects in Table 3.1)*
EVIDENCE	r:[N][51]	R believes the situation denoted by the locution N to be true.
JUSTIFY	r:$F(N_w)$	The perlocution made in N by W is felicitous[52] according to R.
MOTIVATE	r:[N]	R desires the act of *bringing about* what is denoted by N.
CAUSE	r:(([S]⇒[N])∧[N])	R believes [recognises] that the situation in N has been caused by the situation in S.[53]

Table 3.2. Logical translations of relation effects for Causation cluster.

Any discrepancy in wording between the Effects and these consequents is because these are only the *consequents* of W's perlocutionary act. When the antecedent is included, the Effects and the logical translation become closer still. Compare:

Effect: [W intends that, as a result of W's locution of S] R's belief of N is increased.

Logical paraphrase: w:(S_w⇒r:[N])

Translation of logical paraphrase: W believes[54] that W's locution of S will cause R to believe N.[55]

Having established the logical paraphrases of these relations, we can now proceed to use these paraphrases to clear up Mann & Thompson's taxonomy of the Causation cluster. Here, again, is Mann & Thompson's taxonomy:

> EVIDENCE is discriminated from the others on the basis that all of the others address an action, but EVIDENCE does not. JUSTIFICATION is discriminated from the remainder in that it addresses a speech action found within the same text, while the others do not. MOTIVATION addresses a potential action, whereas REASON does not. (Mann & Thompson 1986:62)

---

[51] See Appendix C for an explanation of the logical notation. The most important symbols used are the following:

u:A: u accepts that A; u believes that A.

u:A: u accepts that A should happen; u wills A or does A.

[52] *Felicitous* is defined in accordance with Searle's felicity conditions (Levinson 1983:238): conditions which must be met for a given perlocution to arise. When a locution satisfies these conditions, its perlocution is felicitous. The satellite in a JUSTIFY relation would typically outline how one such felicity condition is satisfied. This is clearest when N contains a performative, as in *I'm Officer Krupke. You're under arrest. I'm Officer Krupke* explains why the performative won't *misfire*, to use Austin's term (Levinson 1983:230): the condition that the speaker be a law enforcement officer is met, so the perlocution is felicitous.

But felicity conditions can also be associated with illocutionary acts of merely stating. Consider again the example from p.1, where (2) is analysed as a justification for (3):

1. The next music day is scheduled for July 21 (Saturday), noon–midnight.
2. I'll post more details later,
3. but this is a good time to reserve the place on your calendar.

The fact that more details will be posted later (S), justifies the statement in N: it makes N felicitous to utter, although the promised "further details" are missing. To be precise, S justifies the fact that N does not follow Grice's Cooperation principle: I won't give you full details now (which is what you would expect of me, by the Cooperation principle), because I promise to provide them later.

If we extend felicity to include 'situationally appropriate', and thence 'following Gricean maxims', then S here has made an infelicitous perlocution felicitous. This is why the example constitutes a JUSTIFY span.

[53] More 'literally': the situation in N is materially implied by S, *and* has actually taken place.

[54] The distinction between 'cause belief' and 'increase belief' could also be captured by a symbolic logic, by using fuzzy logic; but I think the point has been adequately made.

[55] For CAUSE, the logical paraphrase would be w:$((N+S)_W$⇒r:(([S]⇒[N])∧[N])): W believes that W's locution of N *and S* will cause R to believe that S causes N.

What EVIDENCE and CAUSE have in common is that they do not involve an illocutionary (N_w) or deontic ($[\![\underline{N}]\!]$) *action*. JUSTIFY provides the background to an illocution, which is a speech act (action.) MOTIVATION encourages the reader to do something, so it also involves an action. EVIDENCE involves not an action by the reader or the writer, but a belief that the state of affairs in N holds.

As for CAUSE, the causation occurs on a different plane. The text asserts that N and S are causally related, without introducing any interesting interpersonal dealings between W and R, other than the standard Informational intent of W to communicate facts to R.[56] So, while there may be an action involved, it is an action within the text world, not outside it.

So EVIDENCE does not 'address' an action, just as Mann & Thompson claim; but neither does CAUSE. N may still *denote* an action. Mann & Thompson appear to use 'addressed' somewhat loosely, because a VOLITIONAL CAUSE necessarily denotes some action, while EVIDENCE needn't. But it seems inappropriate to claim that somehow

Jenny's not coming. Her mother just called from Tiffany's. (EVIDENCE)

does not 'address' an action, whereas

Jenny's not coming. A mate of mine knows her doorman. (JUSTIFY)

does. Because of this terminological confusion, I don't believe Mann & Thompson's 'addresses an action' criterion is helpful in distinguishing between relations.

The other distinctions Mann & Thompson make fall out more easily from my logical analysis. JUSTIFY is directly concerned with the property of an illocution (speech act) in the nucleus; no other relation includes such metalinguistic reasoning in its effect. That is, the other relations may also be present in illocutions, but none of them consist of a metalinguistic *evaluation* of an illocution. MOTIVATION is concerned with a deontic event (an event which *should* be done); an event which should be done is normally a potential event. As Mann & Thompson argue, no other relation considers a *potential* event.

Finally, EVIDENCE is distinguishable from CAUSE because the consequent of EVIDENCE has epistemic force: S causes belief in N, not N itself. The fact that the car park is full did not *cause* the record crowd at the MCG (in fact, the converse probably applies); but it does cause me to *believe* that there is a record crowd.

3.3.3. Relation pairings.

M&T (1987:9) present a list of rhetorical relations in which certain relations are grouped "according to a specific kind of resemblance. Each group consists of relations that share a number of characteristics and differ in one or two particular attributes." In effect, they've done much of the work of recognizing a taxonomy already; but they haven't introduced an overall scheme, tying all the relations together.

---

[56] So I do not claim there is no causation at the perlocutionary level. Stating a CAUSE span still causes R to believe some statement (which happens to also involve a causation). But because CAUSE is an Informational relation, the consequent at the Interpersonal level is ignored in favour of the causal consequent at the Informational level, which is $[\![N]\!]$.

To put it less formally: if asked what is *caused* in EVIDENCE span

I can't find any parking, so there must be a party next door,

we would answer something like "the situation cause the driver to believe/say there's a party on." This is causation at an Interpersonal level. But faced with a CAUSE span like

I couldn't find any parking, so I left,

we would not say "the situation causes the driver to say that the lack of parking causes her to leave", which is the full Interpersonal causation ($r:([\![S]\!]\Rightarrow[\![N]\!])$, but simply "the situation causes the driver to leave", giving just the Informational consequent $[\![N]\!]$.

Besides the CAUSE cluster, ENABLEMENT is grouped with MOTIVATION; EVIDENCE with JUSTIFY; ANTITHESIS with CONCESSION; CONDITION with OTHERWISE; INTERPRETATION with EVALUATION; RESTATEMENT with SUMMARY; and SEQUENCE with CONTRAST (JOINT is included in this group in the appendix, although JOINT is not considered a relation per se.) I will now attempt to justify some of these groupings in the same way as in 3.3.2.

ENABLEMENT–MOTIVATION.

Using deontic logic as above, we can model ENABLEMENT. The effect of ENABLEMENT is *R's potential ability to perform the action presented in N increases*. We can model this in formal logic as $\Diamond((r:[\![\underline{N}]\!])\wedge[\![N]\!])$: *it is possible that R may bring about the action denoted by N*.[57]

ENABLEMENT and MOTIVATION ($r:[\![\underline{N}]\!]$) are thus both deontic relations: they involve R bringing about some action.

JUSTIFY–EVIDENCE.

The similarity between JUSTIFY and EVIDENCE is intuitively obvious, but somewhat less easy to formalise. Both relations are used in argumentation: EVIDENCE, to substantiate a claim using an inferential chain; JUSTIFY, to support a claim using external authority. If R accepts W's authority, then $(r:F(N_W))+>(r:[\![N]\!])$: if you accept my right to say N, then all other things being equal (conversational implicature), you also accept that N is true. This means that a JUSTIFY Effect conversationally implicates an EVIDENCE Effect. If you believe that I'm Officer Krupke, and that police officers can arrest people (felicity conditions for *You're under arrest),* then you'll also believe that you are, in fact, under arrest.[58]

That justification can behave like evidence should come as no surprise. Indeed, bringing about $r:[\![N]\!]$ (belief in N) is the whole point of W uttering S in the first place, whether in a JUSTIFY or an EVIDENCE span. Both relations are distinct from CAUSE, in that neither can ever imply $r:([\![S]\!]\Rightarrow[\![N]\!]\wedge[\![N]\!]$.) The fact that I'm Officer Krupke does not *cause* you to be under arrest; the fact that I *say* so—does.

The reason why MOTIVATION is distinct from EVIDENCE and JUSTIFY, on the other hand, is that it is not felicitous to say $r:[\![N]\!]$ in a MOTIVATION context: *You should get out more. Staying home all the time makes you a sociopath. Now, you do accept this. ??Therefore, you are getting out more.*[59] This is because the consequent of MOTIVATION is not $r:[\![N]\!]$ but $r:[\![\underline{N}]\!]$: *Therefore, you think you* ***should*** *get out more*.

ANTITHESIS–CONCESSION.

ANTITHESIS and CONCESSION are both Presentational relations; they are both adversative (meaning that both imply the weaker relation of CONTRAST), and both have the same

---

[57]To be more explicit: It is physically possible that: both R wishes to bring about N, and N actually happens. 'Both R wishes N to happen, and N actually happens' is a way of saying 'R volitionally makes N actually happen', i.e. 'R makes N happen; R does N'. The formula thus simply says: 'R can do N' (as a result of the information in S).

[58]To use a non-performative example: if you believe it makes sense for me to say you should mark the Music day in your diary without giving any more details, because I've promised I'll give you the details later (felicity condition), then all other things being equal, you will accept that you should mark the day in your diary.

[59]Which may sound akin to the performative in *Your eyes are getting sleepy*, but does not constitute a valid implicature.

effect (increasing positive regard for N.) The only substantial difference between them is that W regards S and N as compatible in CONCESSION, despite any apparent incompatibility, whereas N and S are perforce incompatible in ANTITHESIS.

So the CONCESSION span *Although the hour was late, they stayed out and partied some more* implies that it's possible both for the hour to be late, and for the revellers to stay out and party—although at first blush, the two are implicated to be incompatible. But in an ANTITHESIS like *We don't want apple juice; we want orange juice*, the speaker claims that wanting apple juice and wanting orange juice are incompatible; only one (N) can hold, and the rejection of S bolsters this claim.

EVALUATION–INTERPRETATION.

M&T themselves provide the reasoning for grouping EVALUATION and INTERPRETATION together:

> The INTERPRETATION and EVALUATION relations involve assessing nuclear material in terms of some frame of reference that is not part of the subject matter of the nucleus itself. The difference is that EVALUATION relates the nuclear situation to a scale of positive regard on the part of the writer, while INTERPRETATION relates the nuclear situation to any other frame of ideas. (M&T:67)

OTHER PAIRINGS.

We have already seen justification for subgrouping the remaining relations, from other taxonomies. RESTATEMENT and SUMMARY are both considered Paraphrases in Longacre's scheme. OTHERWISE is simply the adversative version of CONDITION, as argued by Longacre, Halliday & Hasan, and (indirectly) Knott & Dale. SEQUENCE, CONTRAST and JOINT correspond straightforwardly to Longacre's and Halliday & Hasan's Additive (CONTRAST, JOINT) and Temporal (SEQUENCE) relations.[60]

3.4. The Tilburg taxonomy.

3.4.1. The relations the Tilburg taxonomy includes.

Sanders *et al.* (1992), working at the University of Tilburg in the Netherlands, purport to create a theory of coherence relations that is not only descriptively adequate (which they accept RST is), but also psychologically plausible. They point out (correctly, I believe) that a theory where all relational propositions are cognitively basic, without any further structure to them, is not psychologically convincing. They investigate which primitives might be pertinent for ordering a limited set of classes of rhetorical relations. As for the descriptive adequacy of their taxonomy, they claim that

> the taxonomy proposed here can be extended, using segment-specific features, to arrive at a complete and descriptively adequate set of coherence relations à la [RST] (Sanders *et al.* 1992:5)

The most important contribution they make is identifying the relational criterion, which I discussed extensively in Chapter 2. As I argued, the relational criterion goes a long way toward removing from rhetorical inventories those distinctions which are properly formal-semantic, and not relational at all (such as Volitionality.)

The actual taxonomical distinctions Sanders *et al.* make, however, add nothing to the work of Longacre (1983) or Halliday & Hasan (1976). The four distinctions they make are: whether the relation is additive or causal (a distinction they name *basic operation);*

---

[60] Although the criterion under which they are grouped together in M&T is explicitly syntactic—the fact that they are multi-nuclear, with no syntactically subordinate span encompassed by the relation.

whether its *source of coherence* is semantic or pragmatic (à la Halliday & Hasan's Internal/External distinction); (if the relation is causal) whether the causal antecedent precedes *(basic order)* or follows *(non-basic order)* its consequent; and whether the negation of any segment is part of the basic operation *(polarity.)* Their taxonomy is outlined in Appendix B.3., and is summarised in Fig. 3.3.:

Basic Operation	**Source of Coherence**	**Polarity**	**Relation**
Causal	Semantic	Positive	CAUSE-CONSEQUENCE
Causal	Semantic	Negative	CONTRASTIVE CAUSE-CONSEQUENCE
Causal	Pragmatic	Positive	ARGUMENT-CLAIM INSTRUMENT-GOAL CONDITION-CONSEQUENCE
Causal	Pragmatic	Negative	CONTRASTIVE ARGUMENT-CLAIM
Additive	Semantic	Positive	LIST
Additive	Semantic	Negative	EXCEPTION OPPOSITION
Additive	Pragmatic	Positive	ENUMERATION
Additive	Pragmatic	Negative	CONCESSION

Fig. 3.3. The Tilburg taxonomy of rhetorical relations.

As it turns out, the Tilburg taxonomy has the least coverage of RST relations of the three taxonomies considered so far: it omits 13 of M&T's relations, where Halliday & Hasan's taxonomy omits 10, and Longacre's 6. There are several serious problems with it. To begin with (as already argued in sections 2.1.3 and 3.3.1.), the order of text spans is a purely syntactic issue, and does not distinguish between rhetorical predicates themselves. So it does not belong in a rhetorical theory, however relevant it may be to a taxonomy of connectives.

Secondly, Sanders *et al.*'s account of adversative relations seems naive in its attempt to force the relation into propositional logic.The sentence *Although he didn't have any political experience, he was elected president* does not mean

LACK_POL_EXPERIENCE*(x)*⇒¬ELECT_PRES*(y,x)*,

although this is what Sanders *et al.* claim. With the possible exception of OTHERWISE, all adversative relations, (and not just the additive version of CONCESSION, as Sanders *et al.* (1992:20) argue) deal with defeated implicature. The sentence given actually translates to

(LACK_POL_EXPERIENCE*(x)*+>¬ELECT_PRES*(y,x))*∧ELECT_PRES*(y,x.)*

A third, quite obvious problem is that the taxonomy is only applicable to a subset of rhetorical relations—those Longacre identified as Basic. This is why the Tilburg coverage of relations is so narrow. Longacre defines basic relations as

> more basic to the structure of discourse, while [Elaborative relations…] are considered to be embellishments, i.e. rhetorical devices. Notice that the basic notions of the statement calculus are essentially the logician's conjoining, alternation, and implication plus temporal—which seems to be a necessary addition. (Longacre 1983:80)

So it is no surprise that Sanders *et al.* limit their analysis to these relations: basic relations are the relations easiest to express in propositional logic (without requiring the deontic or epistemic extensions exploited above.) Sanders *et al.*'s taxonomical parameters are such that they can only cope with distinctions expressible in propositional logic.[61]

[61]The distinction between illocutions and locutions, used for the Source of Coherence parameter, does not require a change to a stronger logical framework—unlike the analyses I have given.

3.4.2. The two CONCESSIONS.

A surprising consequence of the Tilburg work is that no less than three different kinds of CONCESSION are postulated within a list of essentially eight relations:

> • a pragmatic causal CONCESSION, e.g. *Although the papers have written about gas accidents several times last year, the risk run by the gas user is much smaller than that of someone in traffic;*
> • a semantic causal CONCESSION, e.g. *In 1969 Hans Hetzel was sentenced to life-long hard labour for murder, although he had stoutly maintained his innocence;*
> • and an additive pragmatic CONCESSION, e.g. *The consumption of mineral water has been advocated strongly over the last few years in the Netherlands, but the results of an investigation in Germany on the composition of bottled water were not so good.*

In the semantic causal relation, a fact in the real world happens despite another. In its pragmatic counterpart, an illocution happens despite the contents of the satellite. And in the additive relation, no causation is present at all; two contrasting illocutions are made, but they are linked by logical conjunction, not implication.

We could concede that the additive pragmatic relation is in fact a CONTRAST (though the question then arises: what is the real difference between this and an additive semantic negative relation?.) But the two remaining relations are still too close for comfort. Both relations' examples pass M&T's criteria for CONCESSION, and so are indistinguishable from the perspective of RST. In fact, one would expect that laypersons would find it difficult to distinguish between the two relations.

This seems to be borne out by Dutch-language[62] experimental work done by Sanders *et al.* When asked to classify relations of text spans according to their scheme, their subjects got semantic CONCESSION right only 27% of the time, and pragmatic CONCESSION 35% of the time. (They got additive CONCESSION right 54% of the time.) The authors argue that this is because both causal relations were frequently misclassified as additive CONCESSION, and suggest that inferred causation interferes with the additive relation. Still, the Source of Coherence parameter was the parameter most often misclassified in the taxonomy (only 65% of all Semantic and 86% of all Pragmatic relations were correctly identified.)

What if language users prove unable to make the distinction between the two types of CONCESSION? The authors assume that, for a taxonomy to be plausible, every pigeonhole in its feature-space must be filled by a distinct taxon—as is the case in classical componential semantic analysis. For instance, one of the arguments Sanders *et al.* give against using temporality as a taxonomical parameter is that temporality is "not [taxonomically] productive": temporal relations can only be semantic, and there would be no relation to fill the space created for pragmatic temporal relations.[63]

But the purpose of a taxonomy should be to capture as many generalisations about the data as possible—not to manufacture artificial distinctions where not intuitively plausible. If people fail to distinguish between the two types of CONCESSION, this is a psycholinguistic fact any taxonomy should recognise and incorporate; the neatness of the taxonomy should not take higher priority than this.

---

[62]The fact that the work was done in Dutch will prove relevant later in the chapter.

[63]Halliday & Hasan (1976) would disagree; they class what Sanders *et al.* consider a pragmatic additive relation, PRESENTATIONAL SEQUENCE, as a pragmatic *temporal* relation. Indeed, they consider this relation the most salient and readily identifiable occurrence of the external parameter (their counterpart to Sanders *et al.*'s Pragmatic Source of Coherence).

So is there any psychological reality to the distinction Sanders *et al.* make between the two types of CONCESSION? Sanders *et al.* claim that the two types *are* marked differently in Dutch; the semantic form with *hoewel*, and the pragmatic with *al*. This is interesting, since English fails to make this distinction—Sanders *et al.*'s claims about *despite the fact that* notwithstanding.[64] Native speakers of Dutch I have consulted have confirmed that *hoewel* and *al* are not intersubstitutable in text, and do have the functions Sanders *et al.* ascribe to them—although *al* is used only in literary Dutch, and there is some confusion about its usage.[65]

However, when the authors presented their subjects with text spans lacking connectives, and asked them to re-insert connectives from their scheme into these texts, their subjects only got CONTRASTIVE CLAIM-ARGUMENT *(al)* right 36% of the time (though they did get CONTRASTIVE CONSEQUENCE-CAUSE *(hoewel)* right 79% of the time.)[66] The errors the subjects made involving *al* are consistent with the confusion about *al* shown by the native speakers I consulted.

Nonetheless, there does appear to be evidence that the distinction between the two CONCESSIONS has psychological reality after all. Still, their conflation in English, and the difficulty Dutch speakers have in distinguishing between them, indicate that the distinction is probably harder to make than for CAUSE and EVIDENCE.

These latter two relations are also conflated at the single-word marker level in English. Despite Sanders *et al.*'s claim that they are rhetorically distinct, the connectives *since* and *because* are intersubstitutable, as Knott & Dale (1993) found. But the distinction between CAUSE and EVIDENCE *is* reflected at the phrasal level *(as a result* versus *it follows that)*, and is also realised by single-word markers in French *(parce que* versus *puisque:* Sweetser (1991:82)) and (as Sanders *et al.* claim) Dutch *(omdat* versus *aangezien.)*

---

[64] Sanders *et al.* claim that *despite the fact that* is only used for semantic CONCESSION, while *although* can be used for both semantic and pragmatic CONCESSION. But both *He isn't ill despite the fact that I saw him at home* and *He isn't ill despite the fact that he'd been skinny-dipping* seem quite acceptable.
Knott & Dale (1993) place *despite the fact* in a subclass of the category occupied by *although*: the former is an EXPECTED PREVENTER OF EVENT, while the latter can be either an EXPECTED PREVENTER OF EVENT or a CONTRAST. However, this distinction is not along the semantic/pragmatic lines Sanders *et al.* argue; it is a distinction between CONCESSION (EXPECTED PREVENTER OF EVENT corresponds to the apparent incompatibility of S and N) and CONTRAST. Knott & Dale (1993:22) themselves argue that the distinction between the two types of CONCESSION claimed by Sanders *et al.* is not marked in English.

[65] Specifically, the two English texts,
He is not sick, ***although*** *he didn't show up for work today,* and
He is not sick, ***although*** *he did go skinny-dipping last night,*
are translated in Dutch as
Hij is niet ziek, ***hoewel/*al/?ook al*** *hij niet op zijn werk was,* and
Hij is niet ziek, ***?hoewel/al/ook al*** *hij 's nachts is gaan zwemmen.*
The native speaker judgements were not consistent, which supports my belief that the distinction is difficult to make. The first speaker consulted (a longtime resident of the U.K.) said that *al* in the first sentence doesn't make sense (although *ook al* is possible "for some people"), and that *hoewel* in the second sentence is nonsensical. The second speaker (a student in the U.S.) commented that using *al* "for this purpose" was becoming obsolete, but *ook al* could only be used in the second sentence.
The third (a student in the Netherlands) thought *hoewel* was acceptable for both sentences; according to him, *al* isn't used in this context, and *ook al* is only acceptable in the second sentence, "but isn't pretty". As further evidence of confusion, this third speaker also said *alhoewel* and *hoewel* had the same function in Dutch; a fourth speaker (also a student in the Netherlands), on the other hand, suggested that, in both texts, only *alhoewel* was appropriate.

[66] The results for CONTRASTIVE CAUSE-CONSEQUENCE, CONTRASTIVE ARGUMENT-CLAIM and (additive) CONCESSION were conflated. Therefore, I cannot analyse the *hoewel/al* distinction for the cases where the nucleus precedes the satellite.

3.6. Hobbs' taxonomy.

Hobbs' (1985) work is somewhat at variance with the other attempts at rhetorical taxonomy investigated in this chapter. He identifies four possible functions for a message in discourse, and uses these functions to classify his 'coherence' relations. The functions are as follows:

> (1) The speaker wants to convey a message. (2) The message is in service of some goal. (3) The speaker must link what he says to what the listener already knows. (4) The speaker should ease the listener's difficulties in comprehension. (Hovy 1985:8)

In the first functional class, the discourse span is coherent because it talks about coherent events in the world. Hobbs finds he cannot explain 'coherence in the world' as either temporal succession or causation: it is a stronger condition than the former, and a weaker condition than the latter. He ends up describing 'coherence in the world' as an *Occasion* relation, where "the first event sets up the occasion for the second." In inferential terms (which Hobbs uses to classify his relations), either "a change of state can be inferred from the assertion of [S] whose final state can be inferred from [N]", or "a change of state can be inferred from the assertion of [N] whose initial state can be inferred from [S]." (Hobbs 1985:10)

An example of an Occasion relation would be *He noticed the broken connection, and took it to his workshop to fix*: "The first clause asserts a change in knowledge that results in the action described in the second clause." The second clause thus constitutes the final state of the narrated situation. Hobbs names CAUSE and ENABLEMENT as special cases of Occasion.

The second functional class is *Evaluation*, and is defined in textual terms: "From [N] infer that [S] is a step in a plan for achieving some goal of the discourse." Examples of this relation are MOTIVATIONS such as *Did you bring your car today? My car is in the garage*, EVALUATIONS like *A funny thing happened to me the other day* (the discourse goal is to be entertaining), or 'metatalk' like *Do you know what I mean?* (the discourse goal is to be understood.)

The third class aims to relate discourse to the listener's prior knowledge, and has two members: BACKGROUND, and EXPLANATION. Hobbs defines EXPLANATION as "infer that the state or event asserted by [S] causes or could cause the state or event asserted by [N]." Thus, *He was in a foul mood; he hadn't slept well that night* is given as an example of EXPLANATION. Hobbs admits that EXPLANATION is merely the inverse relation of CAUSE, and exists as "a reason for telling a story backwards." (Hobbs 1985:13)

The final class of rhetorical relations, the *Expansion* relations, "expand the discourse in place, rather than carrying it forward or filling in background" (Hobbs 1985:14). To unravel the metaphor, expansion relations provide peripheral, non-narrative information about current discourse referents, which takes place in the current, narrated-about discourse world, rather than linking the discourse to other discourse worlds.

In short, they are ELABORATIONS, as a list of relations belonging to this class confirms: PARALLEL;[67] ELABORATION;[68] EXEMPLIFICATION (which RST analyses as a special

---

[67] According to Hobbs, PARALLEL holds when the same predicate can be inferred to hold between the two text spans, as in *He had a fine pair of gloves, and along with the gloves he had a cheap camera*.

[68] Hobbs defines ELABORATION as a special case of PARALLEL, where the same *proposition*—predicate *and* arguments—holds; e.g. *John can open the safe; he knows the combination*, where ◊OPEN*(john,safe)* can be inferred from both clauses.

case of ELABORATION); GENERALISATION; CONTRAST; and VIOLATED EXPECTATION (CONCESSION.)

In a roundabout way, Hobbs' functional distinctions are reminiscent of Halliday's distinction between interpersonal, ideational and textual meanings. Hobbs' Occasion relations talk about relations between states in the real world—so they naturally correspond to Informational relations, incorporating such relations as CAUSE, SEQUENCE, and CONDITION. His Evaluation relations, geared to attain specific 'discourse' goals, are congruent with Interpersonal relations, where the speaker typically seeks to get the listener to do something; my analysis above of relations like EVIDENCE and ENABLEMENT is consistent with this goal-oriented analysis.

Hobbs' Evaluation class is somewhat broader than the Interpersonal class, including 'subjective' elaborations like EVALUATION, and metatextual commentary. RST would normally not consider metatextual material coherent, since it doesn't turn up in their written texts. If RST did consider these metatextual relations, they would probably be classed as EVALUATIONS.

Hobbs' two remaining classes are somewhat more problematic. The discourse goal of integrating current with prior knowledge is a discourse goal independent from any other. Yet it is hard to see how relations serving this goal can form a class distinct from Occasion relations, which show how the world is coherent. The most telling point against the rhetorical independence of Hobbs' BACKGROUND-EXPLANATION class is that it includes EXPLANATION—while CAUSE (to which EXPLANATION is tantamount) is an Occasion relation. This means that what is, in effect, a single rhetorical relation can be placed in two classes because it can satisfy two discourse roles.

But this is not classifying relations as such, but functions of relations. If used to classify *relations*, this double categorisation of CAUSE is too confusing to be useful. In any case, it is too difficult a judgement for an analyst to make. How can we distinguish whether a causation is mentioned as an explanation, or as reportage about the world? If the only distinction is the order of text spans, as Hobbs hints, then the distinction degenerates to the CAUSE/RESULT distinction I have already rejected as merely syntactic; but this would be at odds with Hobbs' own discourse-goal criteria.

The real point of Hobbs' third and fourth classes is to distinguish between two types of Elaborative relations: ELABORATIONS proper, and BACKGROUND. The distinction is useful, but needs to be formulated more carefully, to prevent other relations from interfering with it.

Hobbs' Elaborative class is quite broad, including not only relations traditionally regarded as Elaborative (such as EXEMPLIFICATION), but also Adversative relations like CONTRAST and CONCESSION. In Hobbs' scheme, only a core of non-adversative causal and temporal relations are regarded as 'coherence in the world': *X because Y* is coherence in the world, but *X although Y* is an ELABORATION, its coherence only textual.

I think Hobbs draws too strict a line between the physical world and the mental model of the world, where the incompatibilities defining CONCESSION and the discrepancies defining CONTRAST are active. If CONCESSION is to be classed as Elaborative while CAUSE is an Occasion, what of the CONDITION/OTHERWISE pair? How do they fit across the Occasion/Elaboration divide? Do they both reflect 'coherence in the world', or are they both artifacts of our mental model of the world?

Hobbs' segmentation of rhetorical relations seems to raise more questions than it answers; although its attention to discourse principles is welcome, it does not adequately justify discarding the existing consensus on rhetorical taxonomy.

3.7. A synthesis of taxonomies.

None of the taxonomies I have considered fulfils all the requirements I made of a rhetorical taxonomy in section 3.1.—that it be comprehensive, externally motivated, and feature-based. I now propose the following as a synthesis of all these taxonomies, which satisfies all these requirements.[69]

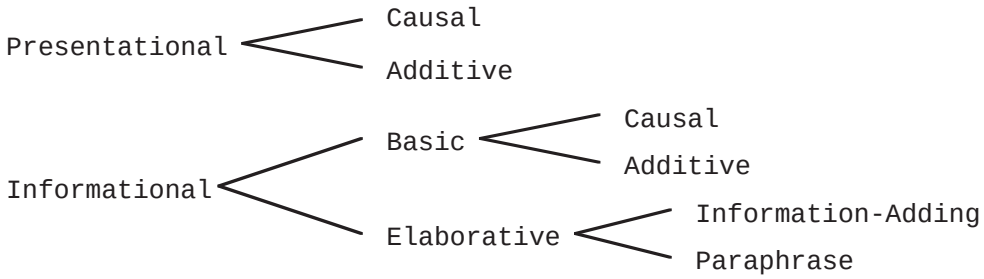

Causal relations:

	Non-Adversative					*Adversative*	
	Modal		*Non-Modal*				
	Deontic	*Non-Deontic*	*Deontic*	*Non-Deontic*			
				Illocut'ry	*Epistemic*		
Pres.	Enabl't		Motiv'n	Justify	Evidence		Concession
Info. Basic	Sol'n.	Condition	Purpose		Cause	Otherwise	
						Modal	*Non-Modal*

Additive relations:

	Non-Adversative	*Adversative*
Present.	*Pres.Sequence*	Antithesis
Inform. (Basic)	Sequence; Joint	Contrast

[69] I obtained a copy of Maier & Hovy's (1993) rhetorical taxonomy too late to integrate it into my discussion; but I discuss it in Appendix D.

Elaborative (Informational):

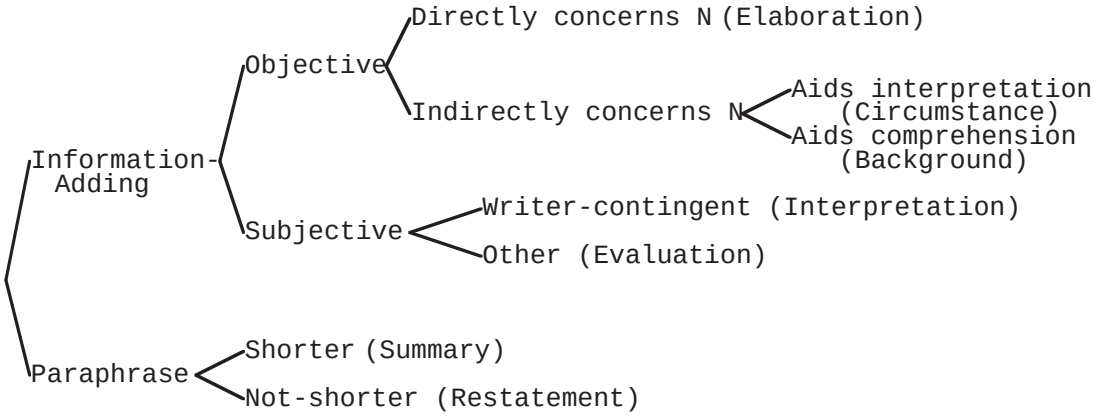

3.7.1. The taxonomical parameters.

The following taxonomical parameters are the features which can be combined in my scheme to describe individual relations.

PRESENTATIONAL/INFORMATIONAL; CAUSAL/ADDITIVE.

In my taxonomy, I retain the M&T distinction between Presentational and Informational relations, and use it in preference to Source of Coherence. This is because of the difficulties arising from this distinction in the Tilburg scheme, and because M&T's perlocution-oriented definitions of relations fit much more neatly into a Presentational/Informational scheme. I also use the Causal/Additive distinction present in Longacre, Halliday & Hasan, and Sanders *et al.*

ADVERSATIVE/NON-ADVERSATIVE.

In addition, I use the Adversative/Non-Adversative distinction implicit in Longacre's and Halliday & Hasan's work. I include CONCESSION, ANTITHESIS, OTHERWISE and CONTRAST as Adversative relations. While there is an intuitive sense that they all go together (they have something to do with negation), a more formal definition of adversativity seems elusive. OTHERWISE, for example, does not involve defeated implicature, nor does it imply CONTRAST (both of which hold for the other three relations.) As a result, Knott & Dale (1993) include *otherwise* not in their Adversative connectives class, but in their Conditionals class.

However, OTHERWISE is similar to ANTITHESIS and CONCESSION (if not CONTRAST), because it involves incompatibility: both OTHERWISE and ANTITHESIS assert that their nuclei are incompatible with their satellites, while CONCESSION has the effect of defeating the presupposition of such an incompatibility.

There is also an intuitive symmetry between CAUSE/CONCESSION, on the one hand, and CONDITION/OTHERWISE on the other: OTHERWISE is a negative counterpart to CONDITION the same way CONCESSION is an adversative counterpart to CAUSE. The following logical paraphrases should make this intuitive symmetry apparent:

OTHERWISE: r:([N]⇒¬[S])	CONDITION: r:([S]⇒[N])
CONCESSION: r:(([S]+>¬[N])∧[N])	CAUSE: r:(([S]⇒[N])∧[N])

Table 3.3. Logical paraphrase of the effects of CAUSE/CONCESSION/CONDITION/OTHERWISE.[70]

Apart from the fact that CONCESSION uses conversational implicature rather than material implication, and that N and S are reversed for OTHERWISE (a discourse effect, and not a difference in the logic of the relation), OTHERWISE is to CONDITION as CONCESSION is to CAUSE, in that both 'adversatives' negate their consequents.

Finally, in logical terms, OTHERWISE is equivalent to a DISJUNCTION,[71] which might be cognitively conflated with CONTRAST (see Sanders *et al.*'s (1992) discussion.) For the purposes of this taxonomy, OTHERWISE will be considered adversative, and the hypothetical equivalent to CONCESSION.

BASIC/ELABORATIVE.

I also introduce Longacre's Basic/Elaborative distinction, considering both as subtypes of Informational relations. This is because the two-way distinction M&T make between Presentational and Informational relations is built into their definitions of RST relations (the former are described as increasing some inclination, whereas the latter—as making R recognise some fact.) However, there is no principled reason why a three-way distinction of Informational/Presentational/Elaborative relations, corresponding to Hallidayan Ideational, Presentational and Textual meaning, couldn't be used instead.

As Longacre pointed out, Elaborative relations have a different underlying logical connection to both Causal and Additive Basic relations (the equivalence connective (≡).) Sanders *et al*. have also stated that Elaborative relations "cut across" their Causal/Additive dichotomy. For this reason, I class both Presentational and Basic relations as either Causal or Additive, but make Elaborative relations insensitive to the distinction; this should not affect descriptive adequacy.

DEONTIC/NON-DEONTIC.

I use the Deontic/Non-Deontic distinction I already made for the Presentational relations in the Causation cluster; I extend it to Basic Causal relations. A relation is deontic if deontic or imperative logic is required to express its effect in formal logic.[72] Having made this distinction, I now outline which relations share it.

---

[70]The logical paraphrases can be translated as follows (RST Effect definitions in parentheses):

OTHERWISE: R believes that N implies not-S. *(R recognizes the dependency relation of prevention between the realization of the situation presented in N and the realization of the situation presented in S.)*

CONDITION: R believes that S implies N. *(R recognizes how the realization of the situation presented in N depends on the realization of the situation presented in S.)*

CONCESSION: R believes that S conversationally implicates not-N, but also that N is in fact true. *(**[Constraint:]** W acknowledges a potential or apparent incompatibility between the situations presented in N and S; W regards the situations presented in N and S as compatible; recognizing the compatibility between the situations presented in N and S increases R's positive regard for the situation presented in N. **[Effect:]** R's positive regard for the situation presented in N is increased (through recognition of the compatibility of situations presented in N and S.)*

CAUSE: R believes that S implies N, and that N is in fact true. *(R recognizes the situation presented in S as a cause for the situation presented in N)*

[71]A disjunction logically corresponds to A∨B. OTHERWISE logically corresponds to A⇒¬B, which by De Morgan's law is logically equivalent to A∨B.

[72]*Imperative* logic encodes the notion of *should*, and commands. In combination with belief logic, it can encode the notion of wanting to do something, and of actually doing something. *Deontic* logic is the equivalent of modal logic for imperative propositions: it deals with what is morally, rather than logically, possible and necessary: in other words, what is permissible or obligatory. See Appendix C, and Gensler (1990).

It seems clear that PURPOSE is the Informational counterpart to the Presentational relation MOTIVATION. MOTIVATION can be rendered as $S_W \Rightarrow r:\underline{[\![N]\!]}$: W's locution of S causes R to desire to do N *(R's desire to perform action presented in N is increased.)* PURPOSE can be rendered as $(m:\underline{[\![S]\!]}) \wedge ((m:[\![S]\!]) \Rightarrow (m:\underline{[\![N]\!]} \wedge [\![N]\!]))$: An agent M[73] desires the state S; as a consequence of this, M does the event in N[74] *(R recognizes that the activity in N is initiated in order to realise S.)*

In both cases, someone wants to do something. In MOTIVATION, it is the reader, as a result of an perlocutionary act. In PURPOSE, the writer is not setting out to make the reader *do* anything; she is simply referring to a state of the world, in which an agent is motivated to do something. The distinction between MOTIVATION and PURPOSE is thus analogous to the distinction already made between EVIDENCE and CAUSE.

Similarly, SOLUTIONHOOD seems to be the Informational counterpart of the Presentational relation ENABLEMENT. ENABLEMENT corresponds to $S_W \Rightarrow \Diamond((r:\underline{[\![N]\!]}) \wedge [\![N]\!])$: W's locution of S makes it physically possible for R to do N *(R's potential ability to perform the action presented in N increases.)* SOLUTIONHOOD corresponds to $[\![N]\!] \Rightarrow \Diamond((m:\underline{[\![S]\!]}) \wedge [\![S]\!])$: N makes it possible for M to do S *(R recognizes the situation presented in N as a solution to the problem stated in S.)* Both relations involve making something possible, but only ENABLEMENT treats the problem to be solved as an interpersonal matter, between the writer and the reader. SOLUTIONHOOD, at least in M&T's definition, confines the 'problem' to the text-world, rather than making it necessarily interpersonal.

The other important distinction between the two relations is that SOLUTIONHOOD involves an explicit request for help (the posing of the problem), which the Solution answers. There is no such request in ENABLEMENT: the enablement does not solve a problem (if it did, it would, by M&T's definition, be a Solution.) Instead, it presents the reader with the means of doing something, before the reader gets around to asking how. If the problem is not posed to the reader, but is embedded in the discourse world, then of course the distinction does not arise.

Thus, *I'm hungry. Let's go to 'Twins.'* poses a problem, and the speaker acts as both questioner and answerer, giving her own solution. *You can get more information by writing to…* is not responding to an explicit problem posed (the problem may still be implicit), but it does allow the reader to do something she may desire to do.

If the distinction still looks tenuous, it is as clear as M&T have made it. An alternative formulation of the two relations would probably mean that SOLUTIONHOOD would have to be classified differently to how I have dome.

MODAL/NON-MODAL.

A further distinction I make between relations is that between Modal and Non-Modal relations. The latter claim that [N], if they are Non-Deontic, and that $\underline{[\![N]\!]}$ *(you should do*

---

[73]That is, an agent in the text-world; not W or R.

[74]This is the *volitional* version of PURPOSE. An example would be *He climbed the fence to get a better view*. The *non-volitional* version of PURPOSE would be notated $O\underline{[\![S]\!]}\phi\iota[\![N]\!]$: N happens because S ought to happen. An example would be *Trees grow tall to catch as much of the sun as possible*. The trees are not volitional agents, but there is a sense where trees catching the sun *ought* to happen, and trees growing tall has to happen to make this possible. Both versions of PURPOSE are therefore deontic relations, according to my criterion.
The philosophy of ascribing purpose in a non-volitional situations seems to me suspect, possibly even anthropomorphic. See M&T's discussion of non-volitional PURPOSE in Appendix A.

N), if Deontic. The former make the weaker claim that ◊⟦N⟧, or ◊⟦N⟧: N is possible, or 'you can do N'. ENABLEMENT, SOLUTIONHOOD and MOTIVATION are all Modal.

I claim that CONDITION is Modal as well. I give the proof for this contention in Appendix E.

ILLOCUTIONARY/EPISTEMIC.

Using the distinctions I have outlined so far, JUSTIFY and EVIDENCE cannot be distinguished. Both relations are Presentational and Causal; neither is Modal, nor are they Deontic. The logical paraphrase of the relations talks about accepting that an illocution is felicitous, versus accepting that a proposition is true; but it seems hard to encapsulate this difference as a differentiating criterion.

Instead, I will exploit Sweetser's (1991) framework for analysing the 'ambiguity' of rhetorical connectives. Sweetser's orientation to this notion of ambiguity is rather different to that of the more formally-oriented researchers we have seen so far. She uses a cognitive-semantic approach, and believes that most notional instances of polysemy can in fact be explained as a metaphorical extension of sense into a new domain, using a well-defined, non-nonce mapping.[75]

When it comes to connectives like *because* (and *so*, *since*, *although*, *therefore*, and other causals), Sweetser believes their ambiguity is not lexical, but pragmatic: "a single semantics is pragmatically applied in different ways according to pragmatic content." Furthermore, she distinguishes not two types of interpretation for causals (as did Rutherford (1970) and his successors), but three: content conjunction, epistemic conjunction, and speech-act conjunction. These types are exemplified by the following three examples:

(1a) John came back, because he loved her.
(1b) John loved her, because he came back.
(1c) What are you doing tonight, because there's a good movie on.

The first type clearly corresponds to M&T's CAUSE, and the second—to their EVIDENCE. What of the third? This particular kind of conjunction is seldom found in written English—presumably because its coherence is somehow harder to derive than for the other two types. So it is doubtful M&T would have anything to say about it at all. But consider the following JUSTIFY span—taken from Mann & Thompson (1986):

I'm Officer Krupke. You're under arrest.

If we lexicalise the conjunction of the clauses, we get:

*You're under arrest, because **I'm** Officer Krupke!*

But this seems to clearly belong with Sweetser's speech-act conjunctions. So Sweetser has arguably set up a three-way CAUSE/EVIDENCE/JUSTIFY distinction.

Sweetser argues strongly against conflating epistemic and speech-act conjunction, as she believes linguists have done in the past:

> (11) He loves me, because he wouldn't have proofread my thesis if he didn't.
>
> […] There is a class of causal-conjunction uses in which the causality is that between premise and conclusion in the speakers mind (as in [11]), and there is another class of uses in which the causality actually involves the speech act itself […] Note that (11) could perfectly well be used to represent a

[75] In other words, there is not a different mapping between core and metaphorical domains for each word. The extension of *I see* from the sensory to the intellectual (=*I understand*), for example, is part of a more general mapping that includes *I feel that, I grasp (a concept)*, etc.

> thought sequence pure and simple; following our omniscient narrator into our heroine's mind, we might find her thinking (11) without any speech act at all. The *because*-clause may in fact secondarily buttress a speech act of assertion (once again assuming [11] to be a speech act), inasmuch as our conversational rules make it incumbent upon us to say things we believe to be true, rather than things we understand to be false. But its primary function, surely, is to explain the epistemic act of drawing the conclusion 'He loves me'. The *because*-clause is fully sufficient as a cause for the act of concluding, but evidence of truth need not by any means be a sufficient cause for the act of stating something. Speech-act conjunction examples […] would more likely refer to the *relevance* or *irrelevance* of a state of affairs as causing or impeding the speaker's action. (Sweetser 1991:80–81)

In effect, Sweetser is saying that EVIDENCE is not an illocutionary conjunction at all. How RST should respond to this challenge is unclear to me. RST theory is expressed in perlocutionary, rather than illocutionary terms, and as far as the *perlocution* of EVIDENCE is concerned, I still believe the relation belongs with the other Presentational relations. Sweetser's classification may mean, however, that in other, more illocution-oriented terms, EVIDENCE may in fact be located half-way between CAUSE and the explicitly illocutionary JUSTIFY.

BACKGROUND.

The main deviation of this taxonomy from M&T's Informational/Presentational division of relations is that I've classed BACKGROUND as an Informational relation. The Effect M&T ascribe to this relation purports to increase an inclination in the reader as much as those of the other Presentational relations. Presumably, M&T classify it as Presentational because they ascribe to W a 'presentational' intent: a BACKGROUND is used to make a text easier to understand.

A more straightforwardly Informational relation like CIRCUMSTANCE, on the other hand, reflects what Hobbs (1985) would call 'coherence in the world': events there tend to co-occur in space or time, as in

> *Probably the most extreme case I witnessed was a few summers ago, when I visited relatives in the Midwest.*

BACKGROUND primarily reflects textual, text-organising coherence: the actual events related may be quite separated in the physical world, and outside the context of the current discourse, an observer may at times be unable to tell they are directly relevant to each other, as in:

> *Someone had left a coffee cup in the office a couple of years back. The office was now colonised with mold.*

Without our encyclopaedic knowledge of the real world, and without the relevance of the two clauses being pointed out to us, we would not establish the connection. For a CIRCUMSTANCE, on the other hand, it can be as easy as identifying temporal simultaneity.

But if BACKGROUND doesn't belong with Informational relations because it involves text-world coherence, then neither do the Informational relations SUMMARY and RESTATEMENT, whose coherence isn't even text-world based, but purely textual. In addition, the actual wording of the Effect of BACKGROUND, *R's ability to comprehend N increases*, doesn't seem to belong with the remainder of Presentational relations. All relations, Presentational and Informational, have a perlocutionary effect, and the perlocution of Background seems characteristically Informational, rather than Presentational: ability to comprehend is what Informational relations are all about. Neither a reader action, nor a reader belief is involved in Background, so it does not have a Presentational type of Interpersonal meaning. For this reason, I have classed BACKGROUND with CIRCUMSTANCE as additional information helping us put the nucleus in context.

The problem resulting from this move is that the difference between BACKGROUND and CIRCUMSTANCE is subtle. The M&T definitions contrast *S increases the ability of R to comprehend an element in N* with *S sets a framework in the subject matter within which R is intended to interpret the situation presented in N*. As was the case with SOLUTION-

HOOD/ ENABLEMENT, the distinction between these relations can only be as clear as M&T's formulation allows. Before any further work can be done on clearing up the distinction between BACKGROUND and CIRCUMSTANCE, a more explicit definition differentiating the two is necessary.

ANTITHESIS.

Classing ANTITHESIS as an Additive relation is a tentative move; but ANTITHESIS is known to be closer to CONTRAST than to CONCESSION. ("ANTITHESIS is a subtype of CONTRAST, as reflected in the definition, while CONCESSION is not" (M&T:11).) Since CONTRAST has been classified as additive, I believe ANTITHESIS (which always implies the weaker CONTRAST) should be considered additive as well. If ANTITHESIS were considered Causal, it would be forced into the same niche in my taxonomy as CONCESSION; I am not sure any clear way to disentangle them could be formulated in that case.

3.7.2. Explaining connectives using the taxonomy.

One of the immediate advantages of this taxonomy is that it can capture some generalisations about the usage of connectives not captured by other approaches, while motivated by independent principles.[76]

For example, *but* can mark only non-modal adversative relations (both CONTRAST and ANTITHESIS claim ⟦N⟧, so they are non-modal):[77]

The long lines outside social security illustrate a lack of jobs, ***but*** *not laziness. (ANTITHESIS)*
[Although Dioxin is toxic to certain animals, evidence is lacking that it has any serious long-term effect on human beings.] Dioxin ***is*** *toxic to certain animals,* ***but*** *evidence is lacking that it has any serious long-term effect on human beings. (CONCESSION)*
Animals heal ***but*** *trees compartmentalise. (CONTRAST)*

But is not applicable to OTHERWISE, which is modal:

[You should have your copy in by Dec 1. Otherwise the existing entry will be used.] You should have your copy in by Dec 1. ?? ***But*** *the existing entry will be used.*

This is consistent with the fact that OTHERWISE is the only adversative relation not to imply CONTRAST.

Similarly, *because* can be accounted for as applying to all non-modal, non-adversative causals:

Bank with us! ***Because*** *we never forget: it's your money. (MOTIVATION)*
There is a competition among trees in certain forest environments to become as tall as possible ***because*** *they want to catch as much of the sun as possible for photosynthesis. (PURPOSE[78])*
You're under arrest, ***because I'm*** *Officer Krupke! (JUSTIFY)*
They must be having a party next door, ***because*** *I couldn't find a parking space. (EVIDENCE)*
Because *I went riding last week, I was sore for three days. (CAUSE)*

Because is not appropriate for modal causals:

[Employees are urged to complete new beneficiary designation forms for retirement benefits if there is a change in marital or family status.] ??Employees are urged to complete new beneficiary

---

[76] Sweetser (1991), discussed above, has a rather different, but extremely interesting approach to the same issue of explaining connective usage.

[77] The examples in this section are taken from either M&T or Mann & Thompson (1986); where necessary to prove a point, rhetorical connectives are inserted into the examples.

[78] Admittedly, this sounds rather more stilted than the other instances.

designation forms for retirement benefits ***because*** *there is a change in marital or family status.*[79] *(CONDITION)*

*??****Because*** *I'm hungry, let's go to the Fuji Gardens. (SOLUTIONHOOD)*

*??****Because*** *I'll give you a free tour of the development, my phone number is 555–9876. (ENABLEMENT)*

However, *because* seems less unacceptable for SOLUTIONHOOD and ENABLEMENT than for CONDITION; and *so* seems capable of dealing with all non-adversative causal relations, except for (prototypically modal) CONDITION:

We never forget: it's your money. ***So*** *bank with us! (MOTIVATION)*

Come for a jog with us, ***so*** *you won't have to stay home and mope. (MOTIVATION)*

There is a competition among trees in certain forest environments to become as tall as possible ***so*** *they can catch as much of the sun as possible for photosynthesis. (PURPOSE)*

I'm Officer Krupke! ***So*** *you're under arrest. (JUSTIFY)*

I couldn't find a parking space, ***so*** *they must be having a party next door. (EVIDENCE)*

I went riding last week, ***so*** *I was sore for three days. (CAUSE)*

*[If I tell him he's a fool, he won't even bat an eyelid] I might tell him he's a fool. ??****So*** *he would not bat an eyelid. (CONDITION)*

I'm hungry. ***So*** *let's go to the Fuji Gardens. (SOLUTIONHOOD)*

Just ***so*** *you can have a free tour of the development: my phone number is 555–9876. (ENABLEMENT*[80]*)*

Why is CONDITION the odd one out? The restriction seems to be that the 'antecedent' in a *so*-clause should be actual, rather than hypothetical. Compare CONDITION to the other two modal relations, ENABLEMENT and SOLUTIONHOOD. It is true that *I am hungry*, and it is true that I am prepared to give you a tour; there is nothing hypothetical about either of these antecedents. But the antecedent for CONDITION is the hypothetical *you may*. Both *so* and (as noted) *because* generate the presupposition that the phrase they precede is not hypothetical. CONDITION seems to be the only relation in which both S and N are hypothetical. For this reason, both markers are unacceptable for CONDITION—but *so* is not unacceptable for SOLUTIONHOOD or ENABLEMENT as well.

Another problem with *so* is whether it introduces the satellite or the nucleus. It normally introduces the consequent, and thus the nucleus;[81] but in PURPOSE, it introduces the satellite. In PURPOSE, *so* seems to have been switched around by intuiting an analogy to CAUSE. S motivates my action N *(I want S,* ***so*** *I do N)*; I do N to attain S; doing N will (I intend) cause S to come true *(N,* ***so that*** *S may be.)* Similar reasoning underlies the use of *so* in ENABLEMENT, but in this case, the satellite of PURPOSE (the aim) corresponds to the nucleus of ENABLEMENT. As the example shows, both the Purposive and the Causal versions of *so* can be used with MOTIVATION.

While *to* cannot even be used grammatically with non-deontic causals, it seems applicable to all deontic causals, not just PURPOSE:

Come for a jog with us, ***to*** *enjoy the evening air.*[82] *(MOTIVATION)*

---

[79] As hinted in the logical paraphrase of CAUSE, *because* here would generate the presupposition that the employees' marital/family status has in fact changed. This would defeat the supposition of a CONDITION relation.

[80] Though this example, too, is slightly stilted, the M&T example of ENABLEMENT, *For a catalog and order form write to…*, works well with a *so* preposed. Either *so* selectionally restricts for imperative satellites in ENABLEMENT, or the *so* is really signalling a MOTIVATION; it's hard to tell which.

[81] There is a reversal of causality between CAUSE *(they're having a party next door, so I couldn't find a parking space)* and EVIDENCE *(I couldn't find a parking space, so they must be having a party next door).* But this is consistent with what Sweetser (1991) would consider the metaphorical extension of causal connectives to EVIDENCE, since what is *caused* is different in the two examples.

[82] The two MOTIVATION examples already used, *We never forget: it's your money. So bank with us!* and *Come for a jog with us, so you won't have to stay home and mope,* cannot be used with *to*—because *to* requires the two clauses

(continued overleaf….)

There is a competition among trees in certain forest environments to become as tall as possible ***to*** *catch as much of the sun as possible for photosynthesis. (PURPOSE)*

To *stop me being hungry, let's go to the Fuji Gardens. (SOLUTIONHOOD)*

To *get a free tour of the development, you can have my phone number: 555–9876. (ENABLEMENT)*

??I'm Officer Krupke ***to*** *arrest you. (JUSTIFY)*

[I couldn't find a parking space, so I must be incompetent] ??I couldn't find a parking space ***to*** *be incompetent. (EVIDENCE)*

??I went riding last week ***to*** *be sore for three days.*[83] *(CAUSE)*

[If I walk into the MCG in the nude, I might get beer cans thrown at me.] ??I might walk into the MCG in the nude, ***to*** *get beer cans thrown at me. (CONDITION)*

However, to get *to* to fit into SOLUTIONHOOD, the sentence had to be rephrased in such a way that it comes close to an ENABLEMENT. Whether this sentence still represents a SOLUTIONHOOD, with constraints peculiar to *to*, or whether it is now an ENABLEMENT, is a question which requires further work—most importantly (as already argued for above), a clearer definition of how ENABLEMENT and SOLUTIONHOOD differ.

As for that old bugbear of rhetorical ambiguity, *and*, OED (1989) lists three senses for *and* as a particle "connecting co-ordinate clauses or sentences":

• Simply connective: a. additive (JOINT): *Love mourn'd long,* ***and*** *sorrow'd after Hope.*

b. adversative (CONTRAST, ANTITHESIS): *He said, I go sir,* ***and*** *went not* (CONTRAST.) *The long lines outside social security illustrates a lack of jobs,* ***and*** *not laziness* (ANTITHESIS.)

• Introducing a consequence: a. the historical sequel or consequence of a fact (SEQUENCE; PRESENTATIONAL SEQUENCE): *You have dismiss'd me,* ***and*** *I go From your breast houseless.*

b. the predicted consequence or fulfilment of a command, or of a hypothesis put imperatively, or elliptically (CONDITION): *Spray with Sanfect* ***and*** *you're safe.*

• Introducing an explanatory, amplificative, or parenthetic clause or phrase (ELABORATION, BACKGROUND, EVALUATION, INTERPRETATION, CIRCUMSTANCE): *The French would certainly have recovered the stolen French provinces whenever they could;* ***and*** *quite right too* (EVALUATION.) *Robert Arrowsmith's father was often thrown into gaol—****and*** *we know what gaols were in those days* (BACKGROUND.) *Yet there is one,* ***and*** *he amongst the foremost* (ELABORATION.) *Steep declines in capital spending commitments pushed the leading composite to a level 0.5% below its high in May 1984.* ***And*** *such a decline is highly unusual at this stage in an expansion* (INTERPRETATION.) *He volunteered to work at the station as a classical music announcer.* ***And*** *that was in 1970* (CIRCUMSTANCE.)

This suggests that all Additive and Information-Adding Elaborative relations can use *and*. (Paraphrase Elaboration would not use it, since A∧A is pragmatically anomalous, and the connective *and* would emphasise that A∧A (rather than A≡A) is being *logically* claimed—undermining the *discourse* justification for uttering paraphrases.)

The use of *and* in CONDITION is an exception to my posited distribution, and deserves further investigation. In this regard, Sweetser's (1991) work looks particularly promising; although she does not address the conditional use of *and* explicitly, she addresses the metaphorical extension of temporal succession to 'epistemic succession' in texts like

Question: Why don't you want me to take basketweaving again this quarter?

Answer: Well, Mary got an MA in basketweaving, and she joined a religious cult. (…so you might go the same way if you take basketweaving.)

> The ordering of [the clauses conjoined by and in the example above] is iconic on the logical processes, rather than on the real-world events involved. (Sweetser 1991:88)

---

it links to have co-referential subjects, and because it does not tolerate extensive syntactic embedding *(??to not have to stay home).*

[83]However, *I went riding last week* ***only*** *to be sore for three days* is acceptable. This may be a metaphorical extension of *to*.

Such a cognitively-oriented approach will probably clear up a lot of ambiguity which more formal approaches cannot properly deal with. Indeed, it may even prove more powerful than the rhetorical-structure–oriented approach I have espoused.

3.8. Computational implications of an RST taxonomy.

Hovy (1991) considers the organisation of RST relations one of the unresolved problems in paragraph planning:

> Practical text structure planning experience accumulated so far indicates that some reorganisation of the RST relations is necessary. Depending on the application domain's needs, certain relations seem naturally to subdivide into sets of similar relations (such as ELABORATION, which gives rise to ELABORATION–PART/WHOLE, ELABORATION–ABSTRACT/INSTANCE, ELABORATION–ATTRIBUTE, etc.), and other relations seem indistinguishable due to our current inability to represent notions such as intention adequately (such as NON-VOLITIONAL-RESULT, VOLITIONAL-RESULT, and PURPOSE.) Fortunately, not all application domains require all relations to be equally finely distinguished: some may require more detailed representations of causality, others may concentrate on object descriptions or historical narrations. In order to accommodate various domain needs, the solution seems to be to structure relations into a hierarchical taxonomy of progressively finer detail. (Hovy 1991:30)

It seems to me important that any such hierarchical refinement occur in a principled way. Moreover, the subdividing features should be clearly identified and scrutinised. For example, the Deontic distinction drawn here seems to cover many of the intentionality problems Hovy encounters (even if we decided against using the Volitionality distinction proper within rhetorical theory.) Identifying which relations it distinguishes should help researchers identify which of those relations will give them difficulties.

A taxonomy should also aid in formulating strategies for resolving any such difficulties, and for disambiguating relations in general. A feature-based analysis affords the text-generator much more flexibility in comparing relations than does a scheme in which each RST relation is atomic.

As an example of such a disambiguation strategy, if we need to stress that a relation is Presentational rather than Informational, we can apply the Source of Coherence parameter,[84] rephrasing the text spans to make the connection explicitly between either locutions or illocutions: e.g. *we can say this because* versus *this is because of the fact that*. In a scheme where each rhetorical relation is an island, we cannot use such generalised parameters to distinguish between relations. Instead, different strategies for disambiguation would have to be worked out for each relation pair—clearly a waste of effort.

I will not be investigating the correlation between intentionality and Informational relations that Moore & Paris (1993:670) use to classify the latter. At first sight, these correlations do not seem to have much to do with my taxonomy (for example, Moore & Paris class CIRCUMSTANCE, CONDITION, CONTRAST, ELABORATION, PURPOSE and SEQUENCE together, as able to express the intention that the reader know about a concept.) It was this incompatibility, after all, that led Moore & Paris to conclude that RST by itself was not adequate to support intentional text planning. Further investigation is obviously warranted—and may eventually lead to a reformulation of RST more sensitive to Speech Act Theory requirements than is currently the case.

---

[84] Even though I rejected this parameter as a taxonomical criterion *in general*, it can still be used to *realise* a Presentational/Informational distinction. While Presentational/Informational is a better way to describe the relations themselves, Source of Coherence seems a better way to describe cohesive links (which is why it is no surprise that it is present in Halliday & Hasan's cohesion-oriented analysis.)

4. The Scott & de Souza programme.

Scott & de Souza (1990) is an influential work in text generation. It is an attempt to exploit RST in text planning, in order to guarantee that the resulting text is rhetorically unambiguous. In other words, the intended rhetorical structure should be recoverable by the reader from the generated text, without error. To bring this about, the text planner is expected to use rhetorical structure information to determine the text's syntax, as well as what textual connectives should be used.

I spend the better part of this chapter scrutinising the heuristics Scott & de Souza come up with. The reason for this is that the authors make several bold (and therefore testable) claims about the relationship between rhetorical structure and text realisation; in particular, that rhetorical structure should be unambiguous, and that connectives should be chosen to effect this. They also make psycholinguistically motivated claims about how text should be structured in order to be easily intelligible.

I believe that ultimately Scott & de Souza's theory is not feasible *in toto*, and that its psycholinguistic foundations are suspect—although many of their proposals are helpful in generating clear and unambiguous text. But looking at why the theory is not feasible, for which aspects of text structure, should prove illuminating—not only in illustrating the limitations text generation is confronted with, but also in examining how and why rhetorical relations are textually signalled in general.

To that end, I look at what Scott & de Souza claim a text generator should do to generate rhetorically unambiguous text. First, I consider the three heuristics underlying their entire theory—that textual markers should be accurate and unambiguous; that rhetorically related propositions should be kept together in a text; and that the text generator should try to make a single sentence out of every rhetorical relation.

Then I consider the practical rules Scott & de Souza formulate, as a result of applying these initial assumptions. Following Halliday's and Longacre's work, Matthiessen & Thompson (1988) have identified three possible ways for clauses to be syntactically related in a text: embedding, paratactic combination ('coordination'), and hypotactic combination ('subordination'.)[85] Scott & de Souza give rules for rhetorically unambiguous text generation involving only the first two alternatives. I evaluate these rules to determine how effective they are.

As it turns out, the rhetorical choices in text involving hypotaxis are much greater than those associated with embedding or parataxis. Because they do not examine hypotactic ways of expressing rhetorical relations, I believe Scott & de Souza make generalisations in their theory which do not hold in all cases. I look at issues arising from expressing rhetorical relations hypotactically, to identify which of their generalisations fail.

Finally, there appear to be several rhetorical relations which cannot be signalled textually at all. This compromises Scott & de Souza's claim that all rhetorical relations should be unambiguously signalled textually. I investigate which relations fall into this class,

---

[85] Using Hallidayan terminology, Matthiessen & Thompson intend the term *parataxis* to describe less interdependent clause combining, including coordination itself, apposition, and quoting. *Hypotaxis* includes non-restrictive relative clauses, reported speech clauses, and the subordinate clauses exemplified by conditionals and causals. The authors reject the term *subordination*, since it fails to distinguish between embedding and hypotaxis.

possible reasons why they tend not to be textually signalled, and what implications this may have for text generation.

4.1. What Scott & de Souza propose: Theoretical heuristics.

4.1.1. Heuristic 1: 'Accurate and unambiguous'.

The first heuristic proposed in S&dS is the underpinning of their entire programme:

> 1. Always generate accurate and unambiguous textual markers of the rhetorical relations that hold between the propositions of the message. (S&dS:50)

As discussed in Section 1.4., the authors present this as the strategy appropriate for text generation, given the impoverished pragmatic resources of a computer interlocutor. Researchers in text generation have often claimed that computer-generated text is likely to be less coherent than human texts, and that text generators need to take particular care to forestall possible resulting misunderstanding. For example:

> In order to enable readers to build a global meaning representation that is equivalent to the writers' original model, they have to be given cues on what the components are, and how to put them together. If the text lacks such instructions, or if the cues are unclear or misleading, then the text is likely to be perceived either as an incoherent set of elements or, worse, as a message having a different goal.[86] (Simonin 1988:146–147)

In addition to these claims, which regard the explicit use of appropriate connectives in a text as desirable, there is also a significant body of psycholinguistic research (summarised in Anderson & Davison (1988:33)), showing that sentences containing explicit textual connectives are better understood by readers than implicitly connected sentences.

Now, it is worth identifying how computer-generated text can actually go awry in practice. Simonin identifies two such ways: the text can be either incoherent; or it can be misleading—coherent, but leading the reader to construe a different rhetorical relation to what was intended.

First, incoherent text. De Souza *et al.* (1989) give the following output of their LETTERA text-generation system as an example of 'stylistically inadequate' text output, which meets at least partly Simonin's description of "an incoherent set of elements":[87]

> *Question*: Was Peter hurt?
>
> *LETTERA*: `Yes, with stabs. The author of the crime is unknown and has fled. The location of the crime is Leme and the date is June 21.` (de Souza *et al.* 1989:224)

They give the following as an example of a 'better' text, which LETTERA should aim for, and in view of which they implemented the heuristics in de Souza *et al.* (1989) (anticipating those in S&dS):

> *LETTERA*: `Yes, with stabs. The author of the crime, which occurred in Leme on June 21st, is unknown and has fled.`

---

[86] Perhaps characteristically of text generation workers, Simonin provides no actual examples of such misleading text.

[87] LETTERA's output is in Brazilian Portugese; the authors provide literal translations into English, which I quote.

The coherence of the first text can still be inferred by a reader *(crime*, and the ellipsis after *date* act as cohesive devices.) But it gives an impression of being 'choppy'. What are normally adjuncts in running prose (location and time) have each been accorded the status of a full phrase, without any obvious functional motivation. The second text rectifies this problem by using embedding. The resulting text sounds a lot 'smoother', and is less obviously computer-generated.

As for misleading text, a message falling under that category would look more like this:

```
Shmuckley is in the building. The system is crashing.
```

Without any further textual or contextual clues, we cannot make head or tail of this. Which is the nucleus in the implicit rhetorical relation—that is, which of the two is the claim that matters? Is the relation between the sentences one of EVIDENCE, JUSTIFICATION, CONCESSION, CAUSE, PURPOSE, SEQUENCE, or JOINT? Depending on which is the case, a whole host of different 'disambiguating' presuppositions can be inferred. Shmuckley could be a klutz of legendary proportions *(satellite–CAUSE–nucleus)*, a notorious saboteur *(nucleus–EVIDENCE–satellite)*, or a helpful computer network guru *(nucleus–CAUSE–satellite.)* His presence could be causing the crash (CAUSE), incidentally preceding the crash (SEQUENCE), or (despite appearances) quite irrelevant to the crash (CONCESSION.)

In normal conversation, these presuppositions would be checked against our background knowledge. We would then select the rhetorical relation that generated the matching presuppositions, to make sense of the text. If we know Shmuckley to be a guru, we will tend to discount the reading in which he causes the crash. Scott & de Souza argue that a computer has too weak a model of the user to be able to assess whether she 'got the message'. For that reason, "their expressions of the message often need to be more explicit than would be ideal." (S&dS:49)

It is true that, with knowledge engineering still in its infancy, the computer model of a user can be ineffectual. How grave a problem this is, though, depends on who is using the system. An expert using an expert system can supply much more background knowledge (and thus requires much less disambiguation on the part of the computer) than a layperson.[88]

But Scott & de Souza fail to distinguish between the two types of text anomaly Simonin discusses. Rhetorically ambiguous text output is quite coherent (after all, several rhetorical (*coherence*) relations *can* be inferred between the text spans.) It corresponds in structure to naturally occurring text, and is contingent on world-knowledge for resolution. The only linguistic principle it violates is Gricean Quantity. Incoherent text, on the other hand, violates principles of information flow and discourse focus. We are left wondering why, for instance, peripheral facts are accorded full sentences, and whether there is some hidden reason for this marked behaviour.

This conflation of the two text anomalies, ambiguity and incoherence, is confusion at a basic level. I believe it is exacerbated by the extension of rhetorical analysis to the intraclausal level, which I commented on in Chapter 2. Researchers are trying to address two levels of sentential realisation at once, using the same machinery. Bringing about coherence in a text—ensuring that there are no clauses in the text which *distract* us from composing a mental model of that text—is an important task for a text planner. RST can ren-

---

[88]However, experts can bring their own problems to reading comprehension. Baker *et al.* (1988:80) report that "the hands-on experience of the experienced NSS technicians actually interfered with their instantiation of facts from the text into their content schema. In several cases, members of this group noted that the text was incomplete or inaccurate (e.g., [the experts would comment] 'it doesn't really work that way.')"

der most potentially incoherent propositions coherent by treating them as ELABORATIONS; and heuristics can be constructed to embed these distracting propositions within other, more focal clauses. Indeed, this is what most of Scott & de Souza's work deals with.

But establishing coherence is quite different work from the other half of the problem: making the coherent text rhetorically unambiguous. Different linguistic strategies are involved in these tasks. Making text rhetorically unambiguous relies on establishing distinctions between textual connectives; establishing coherence finds these distinctions less crucial, and is more prone to exploit more ambiguous syntactic mechanisms, like embedding. Rhetorical non-ambiguity also relies less on discourse focus than does establishing coherence.

In addition, rhetorical non-ambiguity requires the generator to exploit the entire inventory of rhetorical relations; whereas coherence can be established using only three or four relations (just enough to make sure ELABORATIONS and the such are properly integrated into the text.) Coherence for the remaining relations would presumably follow as a by-product of their linguistic realisation, either by explicit textual connectives, or by other cohesive links, like anaphors.

So what is an acceptable heuristic at one level of text generation may not be acceptable at another. The requirements made for unambiguous and for coherent text can conflict. In fact, even within the supraclausal level, where explicit text connectives are most readily available, S&dS's claim is overly optimistic. As I will argue, quite often the appropriate rhetorical markers—if they even exist—simply are not unambiguous.

4.1.2. Heuristic 2: 'Keep the propositions together'.

Scott and de Souza's second heuristic is an echo of Mann & Thompson's Adjacency Principle:

> 2. Keep the propositions of a rhetorical relation together in the text. (S&dS:54)
>
> The text spans of each schema application[89] constitute one text span (M&T:6)

Scott & de Souza justify this heuristic by appealing to the psychology of language comprehension: long-distance dependencies and digressions place a high processing cost on the listener. This places heavy demands on working memory, and slows comprehension. Similarly, undoing temporarily conjectured structures (in the way a typical back-tracking NLP parser does when it finds a discontinuous relation component) is costly in processing terms.

Unsurprisingly, the types of texts considered by M&T keep their propositions together ("a very few texts, typically advertisements in which the title line plays a role in the body of the text, can be analysed only if the adjacency constraint is relaxed" (M&T:8).) They are all written texts, and mainly expository ("administrative memos, magazine articles, advertisements, personal letters, political essays, scientific abstracts and more" (M&T:20).) In such texts, digression and long-distance dependency play a small role; they are more characteristic of spoken discourse, as Fox (1987) has established.

Given the probable social context of computer generation of text (instructional or expository), limiting rhetorical structure through the adjacency principle is unlikely to prove a major hurdle. This is despite the fact that adjacency is known to be a simplistic view of spoken communication. For example, in her analysis of conversational moves (based on spoken dialogues), Reichman (1985) includes an *Interruption move*;

---

[89] That is, the text spans related in a nucleus–satellite or multi-nuclear rhetorical constituent.

> A context space's status slot having such a value reflects that it is expected that the speaker will return and complete discussion of this interrupted context space after the digression has been completed. (Reichman 1985:64)

Many of Reichman's moves correspond to M&T rhetorical relations (*Restatement, Further-Development, Support, Subargument Concession*), the Adjacency Principle in RST is irreconcilable with digression.[90] This is not to say that Reichman's model of discourse is unstructured, or even that that does not incorporate a notion of adjacency. But Reichman's discourse structure is more conceptual, whereas RST makes its structural judgements based directly on the text. In other words, Reichman considers a different level of discourse structure to the rhetorical level M&T consider; this is part of the reason why researchers are reluctant to apply RST to spoken discourse.

But again, given the genres of texts likely to be produced by computer, this constraint on the texts is quite reasonable and uncontroversial.

4.1.3. Heuristic 3: 'Single sentence'.

Research done by Jarvella (1970, 1971) suggests that readers purge their lexical working memory at the end of a sentence. For example, when subjects read the following sentences in running text:

> *Having failed to disprove the charges, Taylor was later fired by the President*

and

> *The document had also blamed him for having failed to disprove the charges. Taylor was later fired by the President,*

Jarvella established that the words *having failed to disprove the charges* are better recalled in the first example, before a sentence boundary intervenes.

Jarvella himself places these results in the context of sentence comprehension:

> Since speech which has been only segmented [into lexemes] may be recalled better than partially interpreted speech, and both might be expected to be remembered more accurately than speech already fully interpreted, the accuracy of immediate verbatim recall may potentially distinguish among several degrees of processing. (Jarvella 1971:410)

Motivated by this research, Scott & de Souza suggest that "rhetorical relations that are expressed within a single sentence are more easily understood than those expressed in more than one sentence", and propose:

> 3. Make a single sentence out of every rhetorical relation. (S&dS:54)

In other words, where possible, realise the rhetorical link between any two text spans *intrasententially*, rather than at a sentence boundary.

This heuristic is, I feel, a misunderstanding of Jarvella's results. The fact that an incomplete (subordinate) phrase is better recalled lexically than a complete sentence does not mean that the former is 'better understood' than the latter. In fact, it suggests quite the opposite. Understanding is something done with the semantic coding of a sentence, not its lexemes. Working memory is purged, not of semantic codings (which Scott & de Souza refer to as "the gist of what was stored"), but of their lexical antecedents.

---

[90] In some cases, embedding relationships in RST will obviate this problem. For example, for a text *A B C*, where *B* is a digression and *C* a resumption, *B* could be considered related to *A*, and *C* related to the *A–B* span, rather than *A*.

This suggests that the incomplete subordinate phrase, better recalled lexically, will likely not yet have been fully processed semantically. We cannot conclude from this that it will be better understood or recalled. Nor can we conclude that the process of mentally linking clauses with rhetorical propositions is any easier within the sentence. There is no inherent difference in textual marking between intrasentential and intersentential constructions; and the semantic content involved in inferring rhetorical relations is not necessarily purged at the end of a sentence, along with its lexical information.

Therefore, Jarvella's work does not prove, as S&dS maintain, that an intrasentential realisation of rhetorical relations is cognitively easier to process than an intersentential realisation. Nor does it disprove it, however. Given that Jarvella's work is not immediately concerned with the building up of text coherence in the mind, we need to turn to other research efforts to get more concrete answers to this question.

4.1.4. Heuristic 3: Van Dijk & Kintsch's model.

Van Dijk & Kintsch's (1983) model of discourse processing is much more helpful in this regard. These researchers posit a model of language comprehension with not only a lexical, but also a propositional short-term memory. In their model, a reader "accumulates semantic units" until a sentence boundary or major clause boundary is encountered. Thereupon, a coherence structure of the text is built up within the reader's mind; and a new processing cycle is started.

They conjecture that, to further assist in building up coherence, some propositions from previous processing cycles are retained for future cycles. In their empirical work, the number of such propositions varies between 1 and 4. Since not all prior propositions can be retained in working memory, propositions are selected according to how high they are located in the hierarchical coherence structure.

Van Dijk & Kintsch's work predates Mann & Thompson's. In particular, it precedes their concept of nuclearity, which Mann & Thompson have claimed is what makes RST distinctive. But van Dijk & Kintsch's model can be phrased in RST terms. The subsidiary information extracted in the processing cycle—the satellites—are discarded from the short-term store once the current cycle concludes. The information left behind is what is used to establish coherence with subsequent text. This information consists of the nuclei of the uppermost schema applications.[91]

---

[91] Van Dijk & Kintsch's theory is very much a propositional theory, rather than a theory of clausal coherence, or for that matter a rhetorical theory. In fact, they decide the issue of which propositions are retained in working memory, not on the basis of rhetorical structure, but on the way propositions share arguments. In the version of the theory used by Fletcher (1981) in his psycholinguistic tests,

> [d]uring each cycle a coherence graph is constructed […] Each proposition in such a graph is connected to the most superordinate proposition with which it shares an argument. If a given proposition shares an argument with two or more propositions at the same level in this hierarchy, it is connected only to the most recent one (the one with the highest [ordinal] number). During the first processing cycle, one of the propositions must be designated as the superordinate. During succeeding cycles new propositions are merely added onto those already in the short-term buffer.
>
> [… To select propositions for retention,] first the superordinate proposition is selected. Next, the most recent proposition is selected from each succeeding level of the coherence graph. If more propositions are required, the most superordinate propositions remaining are selected in order of recency. Any time a selected proposition contains another proposition as an argument, that proposition will be selected next regardless of its position in the coherence graph. This process halts as soon as *s* propositions have been selected. The value of *s* is a parameter of the model which varies from one text to another. (Fletcher 1981:566–567)

But this propositional orientation of Van Dijk & Kintsch's theory does not mean that it is incompatible with RST. Kintsch & van Dijk (1978) argue that a strategy for selecting propositions for retention must be based on both *recency* and *importance*. Their method of determining importance is cohesion-oriented; but I do not see why a more rhetorically-oriented approach would not yield consistent results. The clausal relations ascribed in RST would *usually* be treated as the top level of propositional relations in a Van Dijk & Kintsch propositional hierarchy. See Section 2.2.2. for an example showing this compatibility.

This is consistent with how M&T imply rhetorical structure is built up. In their rhetorical tree structures, each schema application has a single nuclear text associated with it[92]—if its nucleus is not rhetorically atomic, then this nuclear text span is the nuclear text span of the rhetorically compound nucleus. When text spans are rhetorically linked, the texts associated with the juncture are the nuclear text spans. The following diagram illustrates this:

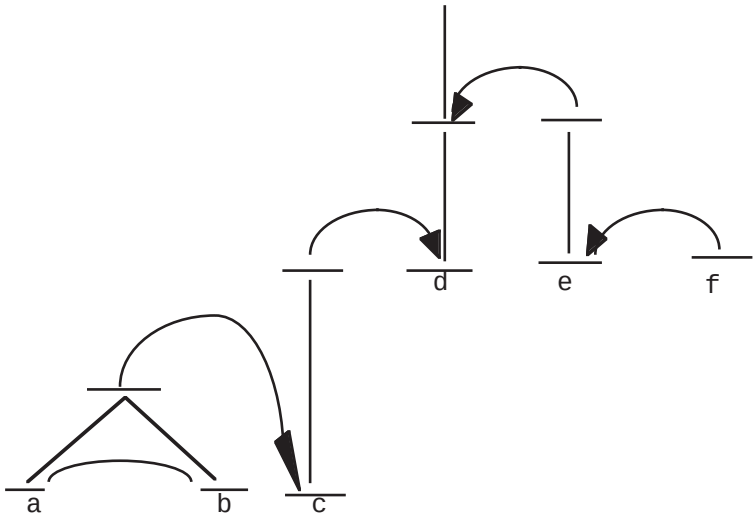

Fig. 4.1. Sample of M&T's tree notation for rhetorical structure.

When span *a–b–c* is linked as a satellite to *d*, M&T's notation treats it as a linkage between two nuclei: *c* and *d*. The same happens when *e–f* is added: nucleus *e* is linked to nucleus *d*. This property of nuclearity is built into M&T's theory:

> […] the nucleus is more central than the satellite in a literal sense. Taking the center as the structural root of the text (the node representing the entirety) and then tracing out from the root, the nucleus is always encountered before the related satellite. Thus the metaphor of centrality is fulfilled. (M&T:38)

In van Dijk & Kintsch's model, if *a–b–c–d* is analysed by a reader in one processing cycle, then only clause *d* would be carried over in working memory to the next cycle, since any further text linkages will be anchored on *d*, and not on *a*, *b*, or *c*.

If van Dijk & Kintsch's model holds, then information need not all be held in one sentence for coherence to be established. Enough information from previous sentences will remain in working memory, that subsequent sentences can be integrated into the mental model of the text without needing to consult long-term semantic memory (which is much less efficient.) However, it may still be the case that coherence is cognitively easier to establish within a sentence, rather than between sentences.

In an attempt to confirm Kintsch & van Dijk's earlier (1978) version of their propositional memory model, without any interference from a distinct *lexical* short-term store, Fletcher (1981) performed memory tests on subjects reading phrases on a screen. The words whose recall was tested could be described according to text position as being in one of four locations. They could be in the most recent cycle; in the next-to-last cycle, but

[92]That is, as long as the relation is nucleus–satellite.

still in the propositional scratch-pad according to Kintsch & van Dijk's criteria;[93] in the next-to-last cycle but outside the scratch-pad; or in prior cycles.

If Kintsch & van Dijk's model is accurate, then propositional short-term memory influences lexical memory, so that non-nuclear lexical material from the penultimate cycle will have been discarded from working memory, and will be recalled no better than material from previous cycles.

Fletcher obtained the following results:

Cycle	*Cued Word Recall (% correct)*	*Word Recognition (% correct)*	*Response time (mean)* (ms)
Last	61	91	1234
Next-to-last: salient	45	70	1385
Next-to-last: non-salient	27	67	1462
Prior	30	68	1478

Table 4.1. Fletcher's (1981) results for lexical recall as a function of cycle of processing and propositional salience ('importance'.)

These results seems to confirm Kintsch & van Dijk's model: non-salient (and, by extension, non-nuclear) material from the immediately previous sentence is as badly recalled as material from prior clauses. They are also consistent with Jarvella's findings, since Jarvella had identified a difference in recall between different clauses in his sentences.

But in addition, lexical recall drops between the current cycle and the penultimate cycle's salient propositions. So while processing *f* in Fig. 4.1., a reader recalls *e* better than *d*, although she recalls *d* much better than *c*, or than *p* in some previous sentence. Does this mean that coherence is more difficult to establish across a sentence boundary? Not necessarily; there may be a threshold of semantic information required for establishing coherence, which is exceeded by the salient propositions. And it should be remembered that lexical encoding is a different domain to semantic encoding; Fletcher's study notwithstanding, it is not clear that a decay in lexical recall implies a commensurate decay in semantic recall.

The psycholinguistic evidence Scott & de Souza invoke to justify their Heuristic 3, then, is not beyond dispute: it does not necessarily prove what they take it to prove. More recent psycholinguistic work, such as Van Dijk & Kintsch's and Fletcher's, seems to lend support to the claim that intrasentential proposition linking is cognitively easier; but this work is not conclusive. It is clear, however, that working memory is well equipped to construe rhetorical relations across sentence boundaries.

4.1.5. Heuristic 3: How many propositions per sentence?

As Scott & de Souza acknowledge, Heuristic 3 causes problems with syntactic complexity. In the worst case, given RST's hierarchical approach to text structure, the entire text would have to be fitted into one sentence. This is obviously untenable: however explicit the rhetorical marking, readers cannot be kept hanging on for the conclusion to a sentence page after page.

Scott & de Souza do not feel they can provide an explicit metric for when a sentence becomes 'too long', although in their view it has more to do with syntactic complexity than length:

---

[93]Namely, high in 'importance'—which I believe is correlated with rhetorical nuclearity, as with *d* in Fig. 4.1., although Fletcher actually used argument-sharing of propositions to determine importance, as already noted.

> the answer seems to lie in some complex combination of factors which include number of words, number of relations, number of propositions, and syntax; factors such as the 'balance' of text also seem to play a role. Just what the magic algorithm is, is unclear to us, and we do not know of any empirical studies on this topic. (There are, however, a number of studies which examine the individual effect of some of these factors) (S&dS:56)

The current consensus amongst readability researchers is that syntactic complexity is a better metric of readability than traditional readability formulae (which tend to rely on word counts per sentence, and on how many of those words are not found in some list of common words.) But as Anderson & Davison (1988) point out,

> difficulty of comprehension is not linked in a simple way to complex features of sentence syntax. That is, complex features of sentence structure do not necessarily present a problem every time they occur. For example, if the context fits the complex structure and justifies its use, the structure may not be difficult to comprehend […] [D]ifficulty of sentence structure is not an absolute value, and depends on interactions with other text features, and with features of the reader. (Anderson & Davison 1988:34)

This implies that, to generate truly readable prose, a text generator must incorporate the very user model Scott & de Souza say is too weak to be relied on. Yet suppose we somehow factor out the reader model from the readability metric, and place some hard limits on embedding complexity (as Scott & de Souza in fact do.) Is Heuristic 3 still a sound way to go about generating text?

Graesser *et al.* (1980) have established that the number of propositions in a sentence is correlated with reading time. Furthermore, because all propositions in a sentence have to be kept in working memory until their hierarchical structure is sorted out, more propositions probably entails more memory effort, until the sentence becomes 'too long'—working memory store is exceeded.[94]

On the other hand, as discussed, an intrasentential realisation of a rhetorical predicate may well be easier to apprehend than an intersentential realisation—particularly if this realisation is hypotactic. This is consistent with Matthiessen & Thompson's (1988) view that hypotactic relations are functionally fitted to the expression of Nucleus–Satellite relations.

So we have two conflicting factors: one motivates less propositions per sentence (processing constraints), while the other motivates more propositions per sentence (ease of apprehending coherence structure.) As a result of this conflict, the readability of a text probably peaks for a certain density of propositions per sentence in that text (because of the second factor), and then falls off (because of the first.) The desirable number of propositions per sentence would be a complex function of the rhetorical relations involved. Intraphrasal propositions, such as relative clauses and adjectives, correspond to Elaborative relations; obviously, more Elaborative relations can be crammed into a sentence than Causal relations, which can only be expressed between phrases or sentences.

Interestingly enough, de Souza *et al.* (1989) list a number of heuristics addressing the issue of how many propositions to pack into a sentence—an issue ignored by that paper's successor, S&dS. The following heuristic from their earlier paper is of particular interest:

---

[94] A sentence becoming 'too long' in this way is not necessarily a disaster, since working memory would presumably behave in the same way as at a sentence boundary; enough salient material would be retained for coherence to still be established. There are limits, however, to how often such a complex sentence can be lexically purged before it becomes too much. We have all had experiences of seemingly unending sentences, where such complexity forces us to use rather different interpretive strategies, ranging from rereading the text to jotting down a phrase marker.

8. Sentences should contain no more than 3 clauses, including embedded ones. (de Souza *et al.* 1989:231)

This heuristic speaks of clauses rather than propositions, which makes it harder to integrate with my propositionally-oriented discussion. Nevertheless, de Souza *et al.*'s rule of thumb (which they admit is "the result of applying native intuitions on the effect of sentence complexity on style") illustrates that the interaction of sentence complexity and readability is more involved than S&dS's Heuristic 3 indicates.

4.1.6. Heuristic 3: Clearer text?

S&dS also argue that Heuristic 3 leads to clearer text, because most rhetorical markers can only be used intrasententially—so, by exploiting them as much as possible, the hierarchical structure of the text will be as explicit as possible, its rhetorical structure delimited by sentence boundaries.[95] Neither of these arguments seems credible in general terms—although there are obvious advantages to applying the heuristic to tidy up ELABORATIONS, as I will discuss in Section 4.2.

As far as intersentential rhetorical linking is concerned, Halliday & Hasan's (1976) analysis of coordination, limited to intersentential links, is a convincing rebuttal to Scott & de Souza. Halliday & Hasan show that phrasal anaphors in constructs like *because of this*, as well as intersentential connectives proper such as *in other words*, are capable of signalling a wide variety of rhetorical relations. In fact, Mann & Thompson (1986:71) explicitly sanction using both these types of connective.

Intersentential rhetorical linking is possible even with prototypically intrasentential relations, like CONDITION. This relation is intuitively associated with the word *if*; and *if*-clauses always realise CONDITION spans within a single sentence. But as Rösner & Stede (1992:210) found, when the CONDITION satellite is complex, the intersentential markers *in this case* or *should this occur* are more appropriate for realising the text than cramming the entire complex CONDITION span in one sentence.

As for making rhetorical structure explicit by exploiting intrasentential marking: only three (possibly four) levels of rhetorical hierarchy seem representable by distinct syntactic mechanisms. The possible syntactic mechanisms are:

- intraphrasal—namely embedding, such as ***For better music,*** *tune to GOLD FM*, and *The* ***young*** *prince arrived.* (Embedding of sentential complements might be considered a distinct rhetorical level; e.g. *The prince,* ***who had just finished reading the thesis,*** *emerged from the library bewildered.)*
- intrasentential (either hypotactic, e.g. *The prince was irate* ***because he had found the thesis boring***, or phrase coordination, e.g. *The prince,* ***who had just finished reading the thesis and found nothing about it interesting,*** *walked out of the library in a fury.)*
- intersentential, e.g. *The prince was irate.* ***This was because*** *the thesis was absolutely dreadful.*

So if we are to rely on syntactic mechanisms alone to distinguish between different levels of a rhetorical tree structure, the tree can have at most four levels (and can span at most 2^4=16 clauses.) If it has any more levels, then at least two levels of the tree will be conflated syntactically. For example, if a tree has five levels, its top two levels will both have their constituents separated by sentence boundaries. But this will mean that there is no *syntactic* way of distinguishing between those two levels. This means that S&dS's approach does not help their avowed intention of making rhetorical structure unambiguously recoverable from text.

---

[95] Of course, S&dS do assume sentence boundaries are still present in the text, since they cannot take Heuristic 3 to its logical conclusion in practice.

As I have argued in Chapter 2, embedding is only available as a syntactic mechanism for Elaborative relations, and Causal relations cannot be readily encoded by embedding. This restricts possible tree structures even further: if the hierarchical rhetorical structure of a text is to be recovered from syntax, not only must there be at most four levels of relations, but the bottom-most relations can only be Elaborative.

The limited syntactic resources for hierarchical marking mean that, to achieve Scott & de Souza's aim of rhetorically unambiguous text, Heuristic 3 is not sufficient. The relations may be expressed adequately by connectives; but a reader would find it difficult to determine which spans are the satellites of which nuclei, if Heuristic 3 were carried to its logical conclusion, and the entire text were one sentence (ignoring the obvious processing difficulties that would arise.)

In view of these syntactic restrictions, something like de Souza *et al.*'s (1989) Heuristic 8 seems a more realistic way of guaranteeing that hierarchical structure is conveyed by syntax as accurately as possible. If there are any more than 3 clauses in a sentence (unless most of them are Elaborative) the generator is likely to run out of ways to distinguish syntactically the hierarchical levels of the clauses with respect to each other.

To conclude: beyond the criticisms that can be laid at the psycholinguistic foundations of Heuristic 3, the heuristic also gives rise to problems with syntactic complexity. In itself, syntactic complexity does not necessitate that a text is difficult to read. But the heuristic fails S&dS's own criterion, that the text should unambiguously reflect rhetorical structure.

To investigate the applicability of S&dS's heuristic, I have identified the different syntactic levels at which rhetorical relations can be realised (or, to phrase it in propositional terms, the levels at which propositions can be conjoined.) I have established that there are at most four such levels. This implies that S&dS's criterion is unattainable for any text with more than four levels of rhetorical hierarchy: that is as many levels as syntax can represent.[96] And not all rhetorical relations are expressible at all syntactic levels. So rather than aiming to realise rhetorical links intrasententially wherever possible—as Heuristic 3 requires—it seems more realistic to limit the number of clauses per sentence, as with de Souza *et al.*'s Heuristic 8.

Scott & de Souza's other rationale for Heuristic 3—that intrasentential markers are more powerful ways to express rhetorical relations than intersentential markers—is also not valid, as shown by Halliday & Hasan's research into sentence connection.

4.2. What Scott & de Souza accomplish: Embedding.

In accordance with their first heuristic, S&dS segment the rhetorical inventory in M&T (to which they add ALTERNATIVE to express disjunction) into classes of relations, to be unambiguously distinguished from each other by each of the three possible syntactic encodings discussed in Matthiessen & Thompson (1988): embedding, hypotaxis, and parataxis.

S&dS base their analysis on the contention, already made by Matthiessen & Thompson, that particular syntactic encodings correspond to particular types of rhetorical relations. They conclude that

---

[96]This point depends on the assumption that textual rhetorical markers are correlated with syntactic level, and that no change in lexical marker alone—leaving the syntactic level unaffected—can change the rhetorical-structure depth at which the relation is realised. This seems to me a valid generalisation, although it would be an interesting issue to confirm by further research.

> certain types of complex sentences are likely to be better expressions of a given rhetorical relation than others, and that the wrong choice of sentence type may lead to the wrong interpretation of the underlying rhetorical relations. (S&dS:56)

This means that syntactic structure, as well as more 'explicit' devices like textual connectives, can serve to signal to the reader the relation or class of relation holding within the current text span.

These different syntactic encodings should be considered *prototypical* strategies, rather than obligatory. Just as the prototypical marker of CONDITION is *if*, the prototypical encoding of CONDITION is a hypotactic coordination, where a nuclear consequent is distinguished from a (semantically and syntactically subordinate) satellite antecedent.

But as we have seen, not all CONDITIONS are encoded in this way. Intersentential connectives, like Halliday & Hasan's (1976) *in that case*, or the conditional *and* in *Talk to him nicely, **and** you'll get a pass*, do not follow this pattern. There are arguably good discourse motivations for not following the pattern in such cases.

In other words, while Mann & Thompson would probably agree with Scott & de Souza's syntactic classification of relations—in fact, they themselves already identified multi-nuclear (paratactic) relations as a distinct category—there are instances where the dominant syntactic encoding strategies are not followed.[97] To fully emulate the syntactic repertoire of humans, text planners should be sensitive to the conditions under which these deviations occur, and exploit them when needed.

On the other hand, we should bear in mind that a text planner need not fully emulate human command of syntax to produce intelligible, acceptable text. This means that, although there are good reasons for humans not to always follow the usual syntactic strategy for nucleus–satellite relations, in practice they avoid the dominant strategies seldom—5% to 10% of the time. If a computer follows the dominant strategies 100% of the time in text generation, the discrepancy will probably not make for unacceptable text, although the reader may notice a certain syntactic inflexibility.

S&dS only consider paratactic and embedding constructions. They consider these constructions as unambiguous ways of conveying the type of rhetorical relation. In their work, they cover only the five relations they believe are unambiguously signalled by such constructions: ELABORATION, ALTERNATIVE, SEQUENCE, CONTRAST, and JOINT. I now describe and evaluate the practical rules they recommend for expressing these relations in text generation.

4.2.1. Textual marking for ELABORATIONS.

As outlined in Chapter 2, embedding involves texts problematic to analyse in a rhetorical theory, since we cannot be sure how deep into the grammar of the clause such a theory should pry. Clearly, though, the texts

> *The prince arrived. He is young;*
> *The prince, who was young, arrived;*

and

> *The young prince arrived*

---

[97] Matthiessen & Thompson (1988:308) give corpus statistics according to which 8% of all Nucleus–Satellite relations were realised with paratactic rather than hypotactic constructs, and 11% of all multi-nuclear relations were realised as hypotactic rather than paratactic constructs. When Elaborating relations were excluded from the Nucleus–Satellite count, the frequency of paratactic Nucleus–Satellite relations fell to 5% (two instances in their corpus).

all have the same truth-conditional meaning. Furthermore, there is a consensus, at least in the computational linguistics community (S&dS, Hovy (1990, 1991), Krifka-Dobeš & Novak (1993)), that they are also rhetorically identical. The two propositions contained within each text are considered to be linked by the same rhetorical predicate, ELABORATION.

The similarity in meaning is too clear not to be captured by a text planner. S&dSo's next two heuristics set up just such an equivalence:

> 4. Embedding can only be applied to the ELABORATION relation. [...]
>
> 5. When embedding, the nucleus of the relation must form the matrix of the sentence, and the satellite the embedded clause. (S&dS:57)

Scott & de Souza justify restricting embedding to ELABORATION by saying it is the only relation for which ELABORATION seems appropriate: it is

> the only valid means by which the propositions of an ELABORATION relation can be combined to form a complex sentence. It is also the only available textual marker of ELABORATION. (S&dS:57)[98]

They mention *by the way* and *to be specific* as possible connective markers of ELABORATION, but conclude these are more likely to be repair markers.

My own feeling is that only some subtypes of ELABORATION can be signalled by certain markers, and even these are often not appropriate. For instance, *to be specific* and *for example* each mark only one type of ELABORATION: GENERALISATION–SPECIFIC and ABSTRACT–INSTANCE respectively. A marker like *by the way* can introduce both a digression and an OBJECT–ATTRIBUTE elaboration—but most of the time, its use isn't appropriate anyway: *The prince arrived. ??By the way, he's young. He knocked on the castle door…*

Other researchers have also noted how the textual marking of ELABORATION is problematical. Knott & Dale (1993), for example, perform an exhaustive corpus search of rhetorical connectives, but fail to unearth any single phrases marking all ELABORATIONS. They find this result "particularly surprising", given how important ELABORATION has been in rhetorical analysis.

There are several ways of addressing this question. Knott & Dale imply that subclasses of ELABORATION, such as EXEMPLIFICATION (ABSTRACT: INSTANCE) are coded with distinct markers. This may be evidence that the different types of ELABORATION are in fact different relations.[99] M&T decided not to treat them as different relations. As I have argued, the only valid way of deciding which distinctions are valid in a rhetorical theory is the relational criterion; to Sanders *et al.* (1992), at least, the criterion rules out qualities like ATTRIBUTION are not relational.

We can also address this issue, and in particular the apparent lack of a consistently usable marker for OBJECT–ATTRIBUTE elaboration, by considering embedding to be the prototypical encoding strategy for ELABORATION. In contrast with other relations associated with connectives, ELABORATION is not formulated in clause-conjoining forms (either hypotactic or paratactic) often enough to motivate the evolution of a single hypotactic or paratactic ELABORATION marker. This is also true for subtypes of Elaboration considered separately, like OBJECT–ATTRIBUTE; perhaps less so for ABSTRACTION–INSTANCE.

---

[98] My intuition is that embedding is suited for all information-adding Elaborative relations in my taxonomy, but to argue this is probably more pragmatic nicety than can be justified in this work.

[99] As Mann (1987) actually proposes.

Furthermore, as argued in Section 2.1., ELABORATION can often be inferred just from the referential meaning of the individual text spans. This reduces the pressure for a textual rhetorical marker for the relation to be grammaticalised.

4.2.2. Where should ELABORATIONS be embedded?

Where the ELABORATION nucleus is complex (namely, spanning more than one clause), it is not clear which clause the satellite should be embedded into. To resolve this question, S&dS propose the following heuristic:

> 6. When embedding, the matrix proposition must be the earliest occurring candidate in the immediate nucleus of the to-be-embedded proposition. (S&dS:58)

For example, consider the case where the nucleus of an ELABORATION is the EVIDENCE span *Since my car is a Renault, it must be French*, and the satellite is *My car is new*. This heuristic motivates the construction of sentences like

Since my new car is a Renault, it's French;
Since my car, which is new, is a Renault, it's French;
My new car is French, since it's a Renault;

rather than

??Since my car is a Renault, it, which is new, is French;
??My car is French since it, which is new, is a Renault.

While these results look impressive, the questionable phrases above would be rejected by simpler, sentence-level heuristics, rather than by an explicitly rhetorical heuristic. To wit: anaphors make bad relative clause heads. The fact that this sentence-level phenomenon (which many researchers would explain as a *syntactic* rather than a discourse phenomenon) falls under the scope of an ostensibly *rhetorical* heuristic, shows the risks of conflating the intraphrasal and the interphrasal in a rhetorical theory.

The heuristic is problematic precisely because of this conflation. For example, it claims that *My new car is French, since it's a Renault* is a correct realisation, even if *My car is new* is an ELABORATION of *My car is a Renault*, and not of *My car is French* or of the entire EVIDENCE span. This means that the text planner is prevented from generating *My car is French, because it's a new Renault*—a sentence whose difference from *My new car is French, since it's a Renault* is not merely cosmetic: it gives rise to the presupposition, for instance, that old Renaults used to be manufactured outside France. Indeed, the heuristic says nothing about why the planner should generate *Since my new car is a Renault*, rather than *Since my car is a new Renault*, although the two clearly have a different rhetorical structure.

Now, it seems nonsensical that *My new car is French, since it's a Renault* should be generated, when *new* is an ELABORATION of *My car is a Renault*, rather than *My car is French*. Why would the adjective *new* end up in the *French* clause? Surely the *French* clause should not be a 'candidate' for embedding?

And yet, S&dS explicitly give Fig. 4.2. as an example for how embedding should proceed, where *a*, *b*, and *c* correspond to *My car is French*, *My car is a Renault* and *My car is new*, respectively. R1 is EVIDENCE, R2 is ELABORATION, and the asterisk indicates the matrix proposition:

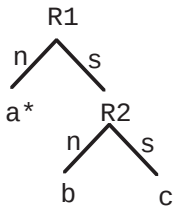

Fig. 4.2. Sample rhetorical structure for embedding in S&dS.

This is another instance in which the earlier de Souza *et al.* (1989) paper makes a less bold and more plausible claim. That paper contains the following two rules:

> *Rule 1:* A Satellite can only be embedded in its Nucleus.
>
> The restriction on which of the Nuclei of a schema can be a candidate home for an embedded Satellite ensures that embedding does not disturb the hierarchical relationships of the RST structure. […]
>
> *Rule 3:* Embedding can occur within the left-most nuclear clause in the structure bearing the same focus value as the candidate clause. (de Souza *et al.* 1989:229)

In other words, a clause can be embedded within another clause at the same or lower rhetorical hierarchical level (*c* in the diagram can be embedded in *b* or a leftmost child of *b*), but not above its rhetorical level (as *a* is.) What seems to have happened is that S&dS discovered the unacceptability of *??Since my car is a Renault, it, which is new, is French*—a purely syntactic effect—and decided to account for it in rhetorical structure terms instead.

To prove that *??Since my car is a Renault, it, which is new, is French* is unacceptable because of an anaphor rule, and not a generic rhetorical principle, we need only replace the anaphor in the embedded proposition with a cataphoric expression. The heuristic would predict

> *??When you see him$_i$ who has performed before the crowned heads of Europe, you'll undoubtedly recognise The Great Houdini$_i$,*

and not the correct

> *When you see him$_i$, you'll undoubtedly recognise The Great Houdini$_i$, who has performed before the crowned heads of Europe.*

Admittedly, cataphora is a marginal phenomenon in English—more so, perhaps, from a sentential than from a discourse perspective. It could be argued that both the tendency to avoid making anaphors heads of relative clauses, and the location in the sentence of elaborative embedding in general, are motivated by the same underlying discourse motivation: concentrate elaborative information around the theme of the sentence, leaving the rheme simple. When anaphora are not involved, the heuristic yields more consistent results. For example,

> *He talked to his brother Peter and to his boss before telling Peter he'd got the job,*

where the apposition *his brother* is juxtaposed with the first mention of *Peter*, is preferable to

> *??He talked to Peter$_i$ and to his boss before telling Peter$_i$, his brother, he'd got the job.*

Still, the difficulties Heuristic 6 run into are proof that a rhetorical theory should know its limitations—in this case, it should not attempt to do what syntax is already well equipped to do.

4.2.3. Dangling sentences.

To prevent the ELABORATION satellite from becoming too complex, Scott & de Souza propose the following heuristic:

> 7. Propositions of a List [Joint] relation should not be embedded if doing so would make the number of remaining propositions in the relation equal to 1. (S&dS:58)

This prevents texts like *The prince came. The prince is young and a Trekkie* from transforming into *The young prince came. He is a Trekkie*. Such a text would destroy the structural integrity of the message, since the JOINT relation in *The prince is young and a Trekkie* would no longer be syntactically recoverable. What's more, to Scott & de Souza the resulting text contains a *dangling sentence* ("information that is only weakly relevant to the message is produced as a separate sentence.") They regard dangling sentences as disruptive to text comprehension, since they come across as weakly coherent afterthoughts. They do so to a greater extent than any other textually unsignalled relation, since ELABORATION

> is the weakest of all rhetorical relations, in that its semantic role is simply one of providing 'more detail'. The information contained in its satellite is thus only weakly relevant to the message. (S&dS:60)

This heuristic seems redundant. If the satellite is complex (here, a JOINT), the satellite must either be embedded into its nucleus as a whole, or not embedded at all. It makes no sense to leave half the satellite as a dangling sentence—not because it would come across as an afterthought (failing to embed the JOINT span gives rise to just as 'dangling' a sentence), but because it compromises the structural integrity of the span: it rearranges the RST tree, making the intended structure unrecoverable.

The structural integrity of the tree is guaranteed by Heuristic 2 (where 'keep together' naturally implies 'don't violate the constituency of the span by inserting half of it into another subtree'.) So Heuristic 7 doesn't seem to be saying anything new.

In any case, 'dangling' sentences do turn up in text, without disrupting texts' rhetorical structure. This is apparent, since ELABORATION is an important relation in classical, *interclausal* RST, which does not extend to analysing embedded, intraclausal relations. There are two good reasons why they would turn up. The first, which I will expand on presently, is that some phrases are simply too heavy to be expressed as adverbs, adjectives or relative clauses.

The second is that the speaker may wish to indicate that the elaborative satellite is salient enough to deserve its own sentence. *He's hired a new, blonde secretary* does not have the same communicative force as *He's hired a new secretary. She's a blonde.* The satellite of the latter has a much more active role in the discourse. Its markedness as a 'dangling' sentence leads (as marked phenomena often do) to the reader implicating many more presuppositions about the sexual dynamics of the situation than does its unmarked form.

But it is doubtful that a text planner created any time soon will have the subtlety required to exploit this kind of markedness successfully. It should also be noted that one of the main problems with text generation to date has been a proliferation of such dangling sentences, a reluctance to embed, frequently leading to stylistically unsatisfying, incoherent text output. Moore & Paris (1993) show this is still an issue with the following output from their text planner:

> You should replace *(SETQ X 1)* with *(SETF X 1.) SETQ* can only be used to assign a value to a simple-variable. In contrast, *SETF* can be used to assign a value to any generalized-variable. A generalized-variable is a storage location that can be named by any access function. (Moore & Paris 1993:656)

The last sentence in this output is an ELABORATION (explanation) of the term *generalized-variable*. Because it is unembedded without any obvious rhetorical motivation, the output looks choppy.

Output like this demonstrates that embedding is an underexploited resource in text generation, and Scott & de Souza, having identified this issue in their own work,[100] have done well to draw attention to it by addressing embedding in the rhetorical domain where such incoherences are conceived. But as we have seen, this has brought problems in its wake—mostly as a result of overestimating the power a rhetorical theory can bring to bear in accounting for grammatical facts.

4.2.4. Syntactic complexity of embedding.

An adjectival modifier and a 'dangling' sentence are, of course, only the endpoints of an embedding cline that also incorporates relative clauses, prepositional phrases, and noun appositions. It seems obvious that expressions higher up on the embedding cline ('lighter') are easier to process than those located lower down ('heavier'): *The young prince arrived* looks easier than *The prince, who is young, arrived*.[101] The following heuristic exploits this:

> 8. Syntactically simple expressions of embedding are to be preferred over more complex ones. (S&dS:60)

Thus, using the phrase *the new car* gives more acceptable text than *the car, which is new*. This heuristic is something of a truism; since syntactically simpler expressions seem *in general* easier to comprehend than more complex expressions, there is no reason why this should not hold in embedding as well as elsewhere.

There are cases where the constituent to be embedded is too heavy to be expressed as a syntactic option high on the embedding cline. For example, the phrase *he is asking the teacher* cannot be embedded into *the student is blond* as a participial clause in English: **the asking-the-teacher student is blond.*[102] The embedding can only be realised as a (heavier) relative clause: *The student who is asking the teacher is blond*. The heuristic is thus not an absolute (although heuristics never are absolute, by definition.) The syntactic realisation of embedding must be as simple as possible, within the limits established by the particular language.

In order to avoid some of the more notorious complexity consequences of embedding, Scott & de Souza add the following heuristic:

> 9. Self-embedding [two levels of embedding] is only allowed in cases where the proposition that is the deeper of the two embeddings is expressed as an adjective or adverb. (S&dS:61)

This allows sentences like

the dog [that likes the [black] cat] is sad,

but not

the dog [that likes the cat [that disappeared]] is sad,

or

the dog [that the cat [that the rat saw] chased] died,

---

[100] See the discussion in de Souza *et al.* (1989)—quoted in section 4.1.

[101] It is not as obvious that we can identify a good metric for syntactic complexity. As I mention below, Frazier (1985) is convinced that none of the metrics formulated to date is adequate.

[102] Although it can in German: *der den Lehrer fragend Student ist blond.*

which have much greater processing difficulty. This is a reasonable condition to place on text intended to be processed easily.

Interestingly enough, in their earlier paper, as well as limiting clauses per proposition to 3, de Souza *et al.* (1989) bar sentences from having more than one level of embedding. The acceptability of *that likes the black cat* has led S&dS to refine this heuristic. (The reason traditional analyses of syntactic complexity have not explicitly addressed cases like *that likes the black cat* is, of course, that an adjective like *black* is not regarded as an embedding comparable to a relative clause in traditional syntactic theory.)

In combination, the authors claim, the embedding heuristics I have described can realise text like

> *My friend George received a long letter from his estranged brother Peter, even though he had told Peter never to contact him,*

in preference to text like

> *My friend George received a long letter from his brother Peter, even though he had told Peter, from whom he is estranged, never to contact him,*

which is rather more difficult to parse, or

> *George, who received a long letter from his estranged brother Peter even though he had told Peter never to contact him, is my friend,*

which has the same propositional meaning as the first text, and could be generated from the same text planner input—but which gives a different rhetorical structure altogether.

4.3. What Scott & de Souza accomplish: Parataxis.

Paratactic constructions are widely used to structure text, particularly spoken text.[103] As a result, there is a widespread feeling that paratactically organised text is easier to process than hypotactic text:

> It is notable that in spoken English, where immediate ease of syntactic composition and comprehension is at a premium, coordinate structures are often preferred to equivalent structures of subordination… Further, spoken English, though less complex in structures of subordination, is more inclined than written English to provide the kind of semantic link that can be made by coordination. (Quirk *et al.* 1972:795)[104]

Notwithstanding the claim that they are easier to process, paratactic markers tend to be ambiguous between a number of rhetorical roles. This is particularly the case with *and*, as shown in Chapter 3; but it also occurs with *or* (which can denote CONDITION and

---

[103] As Chafe (1988) reports, 44% of all intonation units in his corpus are linked only by intonation; of the remainder, 50% are linked with *and*. He also argued that connectives considered subordinating in traditional grammar usually don't behave much differently to coordinating connectives in spoken discourse.
Chafe believes a likely reason why subordination is avoided in spoken language, but sought out in written, is that "speakers, as opposed to writers, have little time to devote to making the linkages between intonation units explicit […] For writers, the absence of a directly shared context and the lack of prosodic and gestural resources make it more imperative to be explicit about connections between ideas." (Chafe 1985:23)

[104] Surprisingly, the psycholinguistic evidence for this seemingly obvious claim is thin on the ground. Flores d'Arcais (1978) reports that "there is some, but still very little evidence for the hypothesis that processing the main clause is easier than processing the subordinate clause". While his own experimental work points to such a difference, I do not regard his results as conclusive, and he finds it difficult to distinguish whether this difference is a syntactic or a semantic effect.
While Frazier (1985) reports "there is considerable evidence for a general principle of minimal structure (the Minimal Attachment strategy) that favors flatter structures over binary branching structures" (Frazier 1985:153), she does not find that any of the current linguistic accounts of complexity satisfactorily account for the data.
But if we cannot even establish the weaker claim that subordination leads to processing complexity as a general rule, we also cannot make the stronger claim that parataxis is easier to process than hypotaxis.

OTHERWISE as well as ALTERNATIVE) and *but* (which can denote all adversative relations except for OTHERWISE.)

As I have extensively argued in Chapter 3, there is method to this ambiguity: most relations expressible by an 'ambiguous' marker tend to have a family resemblance.[105] S&dS, however, do not choose to exploit such resemblances in their work (with the exception of the CAUSE cluster.) Instead, they concentrate on preventing any use of parataxis in text from being rhetorically misleading. To do this, they apply a principle already proposed by Matthiessen & Thompson (1988):[106]

> 10. Paratactic Coordination can only be applied to multinuclear relations. (S&dS:65)

The notional 'ambiguity' of these connectives is further constrained by the following heuristic:

> 11. The paratactic marker *and* must only be applied to SEQUENCE and LIST [JOINT], *but* to CONTRAST, and *or* to ALTERNATIVE. (S&dS:67)

The only thing this heuristic accomplishes that Heuristic 10 doesn't is to prevent *and* from encroaching on CONTRAST. However, when *and* is used to signal CONTRAST in natural texts (as in *Guess what? I'm smart* ***and*** *you're not*), the semantic contrast between the two text spans is usually readily inferrable.

Only rarely does the reader not know whether the writer intends the two spans to be in contrast or not. Consider the classic example: *He's a republican, but honest*. This is an example notorious for generating the presupposition *Republicans aren't honest*. Yet if a reader is already familiar with the political bias of the writer, then the information in the presupposition is already available to her. And in that case, she could still infer the contrast even from *He's a republican and honest*.

So in most cases in real text (as opposed to example sentences divorced from context), *but* is probably not essential for establishing CONTRAST. The text span semantics can do that job by themselves. Still, this only holds for most cases, not all; always using *but* instead of *and*, just to make sure, should not make computer output any less acceptable.

In propositional semantic terms, all multi-nuclear relations are symmetrical except for SEQUENCE. *A, then B* does not imply *B, then A*. *A and B*, however, does imply *B and A*.[107] This allows S&dS to formulate the following heuristic:

> 12. Propositions of all relations except SEQUENCE can be reordered during paratactic coordination. (S&S:68)

The pragmatic effects of such reordering are not ignored in the formulation of this heuristic. In fact, they are mentioned explicitly:

> This flexibility is often useful, especially in cases where the order of presentation affects the thematic flow of the text. For example, (1) would be more appropriate than its alternative (2) if the preceding sentence were (3), and vice versa if the preceding sentence were (4).
>
> (1) The printer is broken and the chapter is due tomorrow.
> (2) The chapter is due tomorrow and the printer is broken.

---

[105] Maybe not all, admittedly. For example, the ambiguity between CONDITION and SEQUENCE for *and* may be best captured in a more abstract, metonymy-driven cognitive-linguistic approach, rather than in a family resemblance between the two relations.

[106] See start of Section 4.2.

[107] But see below for *but*.

(3) The printer always fails when I most need it.
(4) I doubt that I'll be able to finish this chapter on time. (S&dS:68; examples renumbered)

Again, the question whether a text generator will have a subtle enough command of its text plan to exploit such nuances successfully. In the example S&dS give, the text generator would have to be extremely subtle, since both (3) and (4) would probably be acceptable to most people after both (1) and (2): the textual coherence is easily strong enough that we can make sense of either ordering.

The pragmatic effects of reordering are most obvious with CONTRAST. Simply substituting *but* for *and* in (1) and (2) above is enough to prevent the phrases from being inter-substitutable after (3) or (4):

The printer always fails when I most need it. The chapter is due tomorrow but the printer is broken/??The printer is broken but the chapter is due tomorrow.

The reason why CONTRAST is not always reversible is presumably that two propositions in CONTRAST are not merely incompatible, but incompatible with respect to the framework set up by preceding text and real-world knowledge. The first item presented would then be the item compatible with that framework, and the second item, preceded by *but*, would be incompatible.

In the example above, the framework distinguishes between what is desirable for an agent, and what interferes with its realisation. The printer breaking down interferes with the desired aim—to hand in the thesis by the due date. The thesis being due tomorrow is not incompatible with this aim; the printer breaking down, *is*. That is why reordering the phrases makes no sense: it suggests that the printer breaking down is desirable, or at any rate not anomalous.

The final heuristic Scott & de Souza give exploits the tendency for shared elements to be ellipted in text span coordination. For example, the conjunction *You will wash the wildebeest. Then, you will wash the aardvark* contains two shared elements, and can be realised simply as *You will wash the wildebeest and then the aardvark.* As with embedding, Scott & de Souza wish to bias their text generator towards generating the most succinct and easy-to-parse text, and they regard ellipted conjunctions as easier to parse than unellipted ones.[108] This means that, if a conjunction giving rise to such ellipsis *can* happen, it *should*. Therefore:

13. The greater the number of shared elements between propositions, the more desirable it is to coordinate them. (S&dS:69)

The actual work of coordinating the resulting phrases, and then ellipting the shared elements, would presumably be done by a subsequent stage in the text generator. Once the computer knows it should realise *wash the wildebeest* and *wash the aardvark* in the same sentence (as this heuristic tells it to), it can go ahead with the task of gapping coordinated elements.

4.3.1. Evaluation.

These are all the heuristics given in S&dS. While they announce an ambitious program of maximally unambiguous, explicit rhetorical marking, the article itself deals only with an analysis of embedding strategies (which often encroaches on the territory of syntactic planning), and a quick outline of paratactic sentence-linking. Now, it is not unreasonable that their heuristics should encroach on syntactic issues—such as relativisation with an

[108]This is a relative judgement; as already argued, so are most judgements about readability. But I do not intend to press this particular point further.

anaphor head, or the syntactic complexity of expressions. A sophisticated text planner needs to be holistic, with levels of linguistic analysis freely communicating with each other in the generation process.

What is disturbing is the possibility that a rhetorical (or rhetorically-motivated) theory could be considered as saying all there needs to be said about text generation—not only at the intersentential level, to which it is most suited, but even at an intraphrasal level. S&dS do not make this claim, but it is implicit in the way they handled the unacceptability of *??since it, which is a Renault, is French.*

No rhetorical theory is this powerful. When we arrive at the embedding level of syntactic encoding, we are in a completely different linguistic environment. The constraints on expression at this level are not necessarily congruent with the constraints at the intersentential level which Classical RST deals with. An adequate account of this linguistic level is perforce more complicated than what any rhetorically-motivated account can deal with.

This is not to say that the intraphrasal level is impervious to rhetorical effects—if we extend the term *rhetorical* to include intraphrasal relations between propositions (like considering relative clauses and adjectives as encodings of ELABORATION.) S&dS's reminder that elaborative expressions typically belong in embedding is something many workers in text generation can benefit from. Likewise, trying to constrain paratactic constructions to multi-nuclear relations is a reasonable venture, given the pragmatic impoverishment of computer interlocutors—although I still suspect the ambiguity problem for rhetorical relations in text generation (as opposed to the incoherence problem) is overstated.

What compromises Scott & de Souza's venture, in my view, is a naive approach to rhetorical ambiguity. The family resemblances of relations I explored in Chapter 3 help to identify precisely how such ambiguity arises, but S&dS have not exploited this resource. When formulating their general principles, S&dS also rely overly on textual markers (as opposed to any other manifestations of coherence—not just cohesive ones) to disambiguate rhetorical relations. As we will see, there are cases where inferencing is the only tool available to sort the rhetorical web out.

This is not to single out Scott & de Souza as being necessarily any more naive about RST than anyone else in computational linguistics involved with it. Indeed, many of the heuristics derived by S&dS are well thought out and appropriate to the task. Rather, it should serve as a cautionary note, now that RST is becoming something of a *lingua franca* of text generation, that computationalists should be careful how they utilise this resource—and not overstretch it.

4.4. Filling in the blanks: Hypotaxis.

To handle the majority of RST relations in accordance with S&dS's criteria—namely, that the textual marking of the rhetorical relation be clear and unambiguous—we cannot use either parataxis or embedding; these have already been used for the handful of relations I investigated above. Instead, we would have to turn to either hypotaxis, or discourse deixis like *because of that*—both of which are beyond the scope of S&dS.[109]

In the more restricted domains of Additive ('multi-nuclear') relations and ELABORATION, S&dS obtain reasonable results, which seem to bear out their theoretical assumptions about how to mark rhetorical structure in a text. Do their assumptions also hold for the remaining 20 relations in M&T? Does the textual marking of these relations

[109] S&dS mentioned that a forthcoming article by Scott, 'A cognitive approach to the generation of hypotaxis' would deal with hypotactic rhetorical expression. I have been unable to locate this article.

give rise to effects which contradict S&dS's assumptions? I will investigate these questions in the next two sections.

A Scott & de Souza-style treatment of these relations would seek a hypotactic, 'clear and unambiguous', distinct connective to mark each of the remaining relations. Several rhetorical relations are associated with what I would call 'prototypical markers' at the interclausal level. By that, I mean these markers most frequently, if not always, signal their relation in particular, and are very strong cues for its presence. The obvious examples are *although* for CONCESSION, *if* for CONDITION, *otherwise* and *else* for OTHERWISE, *because* and *so* for CAUSE and EVIDENCE (but see below), infinitival *to* and *in order to* for PURPOSE, and temporal connectives like *when* and *while*, or spatial connectives like *where*, for CIRCUMSTANCE.

These markers are generally only valid at the intrasentential level.[110] A different repertoire is used at the intersentential level; these are the connectives investigated by Halliday & Hasan (1976). If anything, these connectives are even more explicit in distinguishing between relations. Thus, as Knott & Dale (1993:21) point out, *as a result* always indicates CAUSE, whereas *it follows that* always indicates EVIDENCE. At the intrasentential level, these two are conflated into *because* or *so*, which are ambiguous between the two—although *since* and *thus*, while also ambiguous, would usually be taken as signalling EVIDENCE.[111]

There are also prototypical markers of rhetorical relations where the markers obtain only at the intersentential level—because their relations, too, are only found at that level. These include *in other words* for RESTATEMENT, *in summary* or *in short* for SUMMARY, *according to [frame of reference]* for INTERPRETATION, and *in my opinion/I think/It seems*, and many other evidentials, for EVALUATION.

There are differences in behaviour between relations signalled primarily by intrasentential and by intersentential markers. Intersentential markers, in particular, are omitted from text more often. World-knowledge and inference are often relied upon to establish how the sentences affected cohere with the rest of the text. For example, there is no overt rhetorical marker in M&T's illustration of EVALUATION, *Features like our uniquely sealed jacket and protective hub ring make our discs last longer. It all adds up to better performance and reliability*. In fact, the particular genre this example is taken from (advertising) would not tolerate the use of an explicit evidential marker: it would apply an (inappropriate) qualification of the claim made.

There are also significant difference in behaviour between the markers of Nucleus–Satellite and multi-nuclear relations. In contrast to the latter, with their straightforward *and*, *but* and *next*, there is a much greater choice of connectives available for Nucleus–Satellite

---

[110] As my supervisor pointed out, I myself use *although* at the intersentential level in this thesis! The instances of such markers used in the speech-act conjunctions analysed by Sweetser (1991), like *What are you doing tonight? **Because** there's a great movie on at the Nicoleum*, would also presumably be considered intersentential. See section 3.7.1. for further discussion of such conjunctions.

[111] Although unambiguous intersentential markers like *as a result* and *it follows that* are available to language users, they usually use the more succinct intrasentential connectives instead, and leave it to world-knowledge to make the distinction between these two relations.
A Scott & de Souza approach would have to opt for the intersentential markers as preferred markers of CAUSE and EVIDENCE, because they are unambiguous. However, given how rarely such connectives are used in text (in contrast to *because* and *so)*, a text relying on them would probably suffer from the same stylistic problem as the Moore & Paris '*in contrast'* text discussed later in this section: wordy connectives are over-used, making the text sound heavy and pompous.

relations,[112] with finer distinctions between the meanings of subclasses of these connectives.

Furthermore, nucleus–satellite markers are more sensitive to textual factors—such as the level of the rhetorical tree at which linking occurs (phrasal, sentential, paragraph); the register of the text; and indeed those stylistic factors Scott & de Souza discounted as reducible to cognitive factors. An instance of this is Stede's (1992) rhetorical-tree–level distinction between *in that case* and *if* for marking CONDITION, discussed above. Another example would be the stylistic difference between *notwithstanding* and *even though*, or the need to choose between nominalisations and *to* clauses to express PURPOSE (Vander Linden *et al.* 1992, Stede 1992, Vander Linden 1993).

Moore & Paris' (1993) text output, in sounding 'forced', is a further illustration of the greater register sensitivity required of these markers:

> `You should replace` *`(SETQ X 1)`* `with` *`(SETF X 1)`*`. SETQ can only be used to assign a value to a simple-variable. In contrast,` *`SETF`* `can be used to assign a value to any generalized-variable.`

The text is quite intelligible, but there is something unwieldy about their use of *in contrast*. My intuition is that *in contrast* typically belongs at higher levels of the rhetorical tree, linking more substantial text spans than what it is linking here—two relatively uncomplicated sentences. Using *whereas* or *while* to conjoin the two phrases into one sentence would give more natural-sounding output. The text gives the impression of a connective slot being filled from a menu, without much sensitivity to the linguistic context.

Fortunately, recent work in text generation has become more aware of these properties of textual connectives:

> A third group of relations is sometimes signalled and sometimes not. The choice of such connectives depends on three factors: the semantic content of the linked propositions, the options for expressing them linguistically, and the choices made so far. The last factor points to avoiding the use of the same connective over and over again. (Stede 1992:3)

> Rather than choosing, by hand, a number of expressional forms that are sufficient to disambiguate a short list of functions of the domain, determined ahead of time to be important, a study in text generation should attempt to deal with the whole range of ideational, textual and interpersonal issues that affect rhetorical and grammatical choice in real texts. (Vander Linden 1993:125)

There remains a problem with prototypical markers when it comes to non-ambiguity. While a rhetorical marker is *prototypically* associated with a relation, it can also be used to signal other, related relations. This is the case with *and* (JOINT), *but* (CONTRAST), and *because* (CAUSE); the related relations were outlined in Section 3.7.2. The Scott & de Souza approach has been to constrain the wider-ranging connectives to their prototypical relations, and to enlist other connectives to express the related relations. This is the whole point of their Heuristics 10 and 11.

Now, this strategy is effective as far as the *related* relations are concerned—such as SEQUENCE (related to JOINT), EVIDENCE (related to CAUSE), or CONCESSION (related to CONTRAST.) The strategy gives *these* relations an unambiguous textual signal—although the stylistic 'fine-print' may mean the choice of connective the strategy makes is sub-opti-

---

[112] For example, in Knott & Dale's (1993) list of connectives, there are seven connectives which could be used to express JOINT: *and, likewise, similarly, also, too, as well, correspondingly*. There are 14 markers of CONTRAST, and presumably just 1 marker or DISJUNCTION (*or*) (but 33 markers of SEQUENCE). On the other side, there are 5 markers for OTHERWISE, 16 markers for CAUSE, 13 for EVIDENCE (most of these overlap), 12 for CONDITION, and 22 each for CIRCUMSTANCE and ANTITHESIS/CONCESSION (Knott & Dale do not distinguish between the two; see discussion below).

mal for other reasons. But the prototypically signalled relation, like CAUSE or JOINT, is still left with an ambiguous marker like *because* or *and*. So rhetorical ambiguity remains an issue for such relations.

4.4.1. Distinguishing adversative relations textually.

This ambiguity is an important problem in distinguishing between the core adversative relations, CONTRAST, ANTITHESIS, and CONCESSION. M&T admits ANTITHESIS is a special case of CONTRAST, so it's not surprising the two relations share markers such as *but*, *however*, and *yet*. But is there any lexical cue to distinguish between the three, or does the distinction rely solely on pragmatic factors?

Let us consider the following examples from Mann & Thompson (1986):

Animals heal, but trees compartmentalise. (CONTRAST)
This book claims to be a guide to all the trees of Indiana. It's so incomplete that it doesn't even have oak trees in it. (ANTITHESIS)
I know you have great credentials. You don't fit the job description because this job requires someone with extensive experience. (CONCESSION)

The M&T definition of ANTITHESIS speaks of an incompatibility between the Nucleus and the Satellite, such that the reader cannot accept both. CONCESSION denies that an apparent incompatibility between the two actually holds. CONTRAST makes no claim as to incompatibility. Now, whether two situations are compatible or not is something to be resolved by world-knowledge and inference. Mann & Thompson's (1986) earlier formulation of THESIS–ANTITHESIS (in an inventory that does not list CONTRAST as a distinct relational predicate) makes the distinction clear by appealing to the writer's mental model, rather than the reader's:

> The 'thesis–antithesis' relationship arises when two conceptions are contrasted, the speaker committing to one and decommitting from the other [… In the 'Indiana trees' example,] the speaker contrasts the idea of the book's being a guide to all Indiana trees, which he or she does not identify with, and its incompleteness, which he or she does identify with. (Mann & Thompson 1986:66)

When viewed in terms of writer's commitment, rather than reader's perception of incompatibility, the ANTITHESIS relation appears to be signalled quite often in text. However, the signal is not primarily the connective *but* (which only indicates that contrast is involved), but an epistemic distancing from the satellite, in phrases like *claims that* or *they want*, or even explicit negatives, as in *We don't want orange juice. We want apple juice.*

In other words, an ANTITHESIS is merely a CONTRAST, with the added information (implicated by the text, though not necessarily by a textual marker) that the writer believes the antithesis, and not the thesis.[113] Therefore, to use M&T's (1987) terminology, the reader should have positive regard for the antithesis, and not the thesis.

Of course, epistemic distance from the satellite does not in itself make up a rhetorical relation; and if that was all that was involved, ANTITHESIS would fail the relational criterion. Mann & Thompson insist, however, that

> both the contrast of content and the contrast of identification between a thesis and an antithesis are combinational effects which do not proceed from the thesis alone. (Mann & Thompson 1986:67)

So does ANTITHESIS have a rhetorical marker of its own? Knott & Dale (1993:30) subdivide adversative relations according to the following intersubstitutability classes in

[113] In my taxonomy in Chapter 3, I address this by splitting CONTRAST between Informational and Presentational: the informational relation is called CONTRAST, and the Presentational relation—ANTITHESIS.

Fig. 4.3. When we insert these markers into our example texts (using one marker from each intersubstitutability class), the following acceptability patterns arise:

*This book claims to be a guide to all the trees of Indiana. (**But/Nevertheless/?By contrast/On the contrary)**, it's so incomplete that it doesn't even have oak trees in it. (ANTITHESIS)*

***(Although/While/Despite the fact that/?Admittedly)** this book claims to be a guide to all the trees of Indiana, it's so incomplete that it doesn't even have oak trees in it. (ANTITHESIS)*

*Animals heal. **(But/?Nevertheless/By contrast/?On the contrary)**, trees compartmentalise. (CONTRAST)*

***(Although/While/?Despite the fact that/?Admittedly)** animals heal, trees compartmentalise. (CONTRAST)*

*I know you have great credentials. **(But/Nevertheless/?By contrast/?On the contrary)** you don't fit the job description because this job requires someone with extensive experience. (CONCESSION)*

***(Although/While/Despite the fact that)** I know you have great credentials, you don't fit the job description because this job requires someone with extensive experience. (CONCESSION)*

***Admittedly**, I know you have great credentials. But you don't fit the job description because this job requires someone with extensive experience. (CONCESSION)*

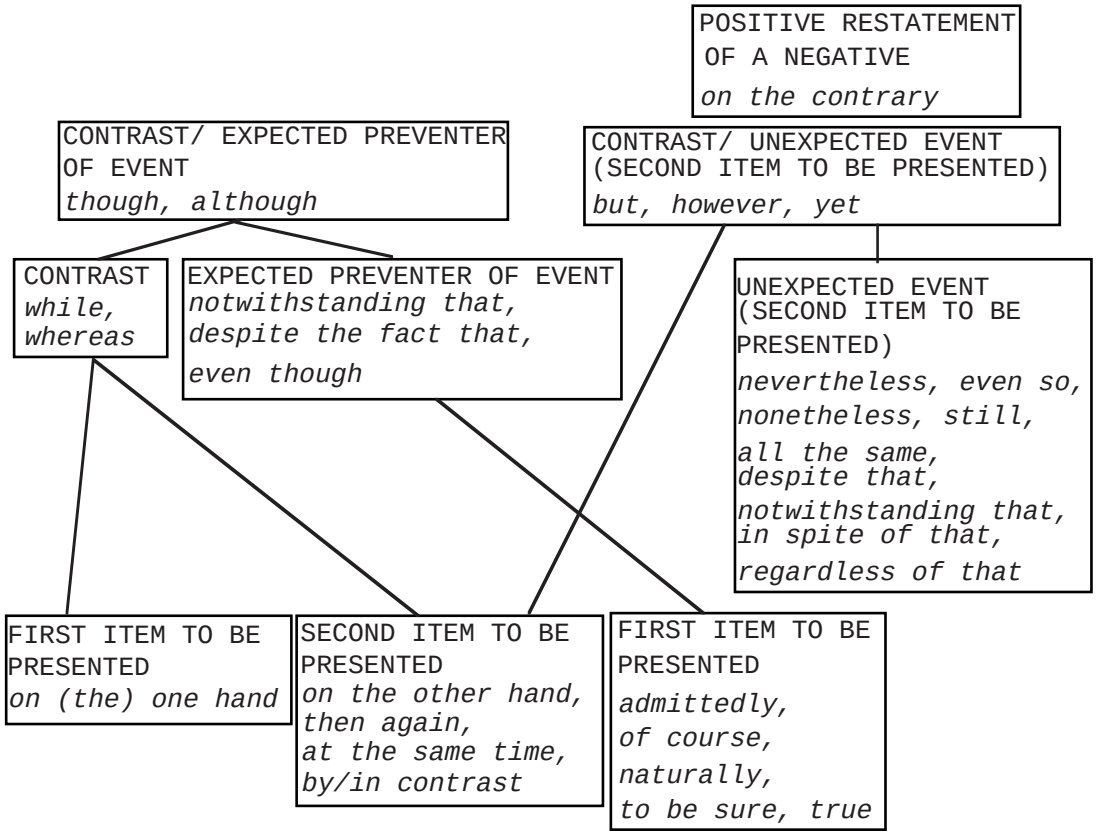

Fig. 4.3. Knott & Dale's (1993) classification of adversative connectives.

The results seem similar to the EVIDENCE/CAUSE signalling problem: the common and short markers tend to be shared among all rhetorical relations. More specific, longer markers are associated with particular relations: *on the contrary* with ANTITHESIS, *by contrast* with CONTRAST, and *admittedly* with CONCESSION. But as we have seen, using these markers alone consistently gives rise to the kind of problem encountered by Moore & Paris (1993) with *in contrast*: the text can sound unnatural.

So what can we conclude about marking prototypically hypotactic relations? We can conclude that the behaviour of their textual markers is more 'fragile' than those of paratactic relations: more context-sentitive, more dependent on register, and displaying greater

variation between intra- and intersentential linking. Because they are more 'fragile', a Scott & de Souza-type approach to rhetorical marking, seeking primarily to minimise ambiguity and disregarding stylistic factors, will be much less successful than for paratactic relations. And it also means that, if computer-generated rhetorical links are to sound natural rather than stilted and menu-driven, the text planner must incorporate a very refined model of style, register, and rhetorical structure.

4.5. Blanking out the fillers: Unsignalled rhetorical relations.

There remain five relations in the RST inventory whose rhetorical marking I have not yet discussed. This is because they belong to the class of relations which, as Stede (1992:3) found, "are typically not signalled by explicit cue words." The relations are: ENABLEMENT, MOTIVATION, JUSTIFY,[114] SOLUTIONHOOD, and BACKGROUND.

The relations fall naturally into three groups. ENABLEMENT, MOTIVATION and JUSTIFY are all Presentational relations;[115] ENABLEMENT, SOLUTIONHOOD and MOTIVATION are all Deontic; BACKGROUND is an Information-Adding Elaborative relation.

4.5.1. Unsignalled Presentational relations.

Some comments can be made on the failure of these relations to be signalled. For instance, Presentational relations in general are not signalled. Presumably, their illocutionary nature makes it easy to infer the coherence of these relations (and the relations themselves) without any need for overt cues.

Of the presentational relations that *are* textually signalled,

- CONCESSION is used as an Informational relation, as well as Presentational. More importantly, the fact that it is both causal and adversative needs to be signalled, to distinguish it from non-causal adversatives like ANTITHESIS, or non-adversative presentationals, like MOTIVATION and BACKGROUND. Causality and adversativity seem to be the two most salient taxonomic parameters for rhetorical relations, so there would be a high functional motivation for signalling these distinctive aspects of CONCESSION, which incorporates both.

- The rhetorical marking of ANTITHESIS and EVIDENCE has already been discussed. In functionalist terms, the markers for these relations have been borrowed from related relations (CONTRAST and CONCESSION for ANTITHESIS; CAUSE for EVIDENCE)—although distinctive subclasses of connectives specific to these relations (e.g. *on the contrary, it follows that*) are available.

- The remaining presentational relations are not distinguished between each other by the salient adversativity or causality parameters; and they are not associated as intimately with a related, textually-signalled rhetorical relation in the way that ANTITHESIS or EVIDENCE are. The functional pressure motivating the creation of Informational markers is Ideationally-based, and primarily intraphrasal. The relations involved with Presentational relations, on the other hand, tend to be intersentential; and the rhetorical predicates involved are more difficult to differentiate using a lexical semantics. Indeed, they are more difficult for *rhetorical analysts* to identify as distinct relations, let alone naive language users. These conditions do not favour the development of explicit, unambiguous textual markers for these vaguer relations.

---

[114]These three relations can all be signalled by *because* or *if*—in a pinch. But this would probably end up misleading the reader, since these markers are more closely associated with CONDITION and CAUSE.

[115]Maier & Hovy (1993), at least, consider SOLUTIONHOOD Presentational as well; see Appendix C.

It should be borne in mind that the difficulty in signalling Presentational relations is of little consequence for computer text generation. With the exception of the illocutionarily-borderline relation EVIDENCE, such relations are likely to be rare in the instructional types of text computers typically produce. And because of the illocutionary salience of these relations, human inferential capacity will usually smooth over any resulting ambiguities.

4.5.2. SOLUTIONHOOD.

SOLUTIONHOOD probably tends to be unsignalled because it is readily inferrable; the manner in which such inferencing occurs is best covered by an explicitly pragmatic approach, such as Relevance Theory (see e.g. Blakemore (1993).) Using the phrase *the solution is* may at times produce acceptable text, as in *I'm hungry; the solution to that is that we go to Fuji Gardens*—so we would not claim SOLUTIONHOOD *cannot* be signalled. But the fact that this text sounds odd indicates that *this* signalling strategy would be dispreferred.

SOLUTIONHOOD covers a large range of satellites. Identifying the satellite as one of the prototypical SOLUTIONHOOD illocutions (questions, requests, descriptions of needs, etc.), and recognising that the nucleus is a solution to this problem, is sufficient to establish that a SOLUTIONHOOD relation holds. And as is the case with ELABORATION, SOLUTIONHOOD covers too varied a set of logical relations to motivate the formation of a single, catch-all connective to signal it. Therefore, once more, the difficulty in signalling SOLUTIONHOOD is of little consequence for text generation.

4.5.3. BACKGROUND.

The distinction between BACKGROUND and ELABORATION is subtle; perhaps overly so. For instance, Mann & Thompson's (1986) example illustrating BACKGROUND, *Hayes just resigned. He's our chancellor*, looks to me like an ELABORATION—particularly since it can be rephrased as *Hayes, our chancellor, just resigned*. It would be interesting to look at intonation phenomena (and their graphological correlates, such as brackets), to see if these (marginally textual) mechanisms make any distinction between the two.

One could argue that the marker *now* can be used to signal BACKGROUND; this makes sense, in that BACKGROUND spans and *now* spans both involve digressions, while ELABORATIONS continue on a topic rather than digress on it.[116] This topicality distinction suffices to distinguish between BACKGROUND and ELABORATION anyway, and narrative structure considerations probably suffice to keeping these relations apart from causals. But, before we can determine whether this is a problem to be addressed in text generation, the distinction between BACKGROUND and ELABORATION needs to be made clearer.

4.6. Conclusion.

A large part of this chapter has been a rather discursive walk-through of the text generation heuristics proposed by one particular group of researchers. In computational terms the chapter has been somewhat monomanic, pursuing all the ramifications of Scott & de Souza's proposals. However, in terms of its linguistic purview, it travels a long and varied path: it begins with computer modelling of the user, and ends with speculations on grammaticalisation. So it would not be amiss to summarise the chapter, and to put it in some perspective.

In Chapters 2 and 3, I investigated problems with the theoretical structure of RST. The computational use of RST has brought attention to these problems, and has made them critical—but it has not created them. Whatever discourse linguists might think of the em-

---

[116] See Hobbs (1985) discussion of the distinction between BACKGROUND and ELABORATION

pirical way computational linguists treat discourse-linguistic constructs, it is not the computationalists' fault that RST ontology and taxonomy have been underspecified.

In this chapter, I turn my attention to how precisely RST is used in a particular text-generation approach. This approach, Scott & de Souza's, makes several linguistic claims, which are perhaps bolder than those of their colleagues. This boldness has helped me focus my attention on some trends in the computationalist use of RST, which I find worrying. In fact, this boldness is a welcome attempt at making strong empirical claims, which can be checked against theoretical linguistic knowledge—rather than non-commital rules-of-thumb, that cause little controversy or excitement. On the other hand, the high regard with which S&dS is held in the computationalist community means that evaluating their positions is relevant not only to their work, but to the entire field.

There are four parts to my investigation of Scott & de Souza's claims. First, I look at the discourse-analytical and psycholinguistic case for and against their theoretical heuristics, which encode the assumptions on which their entire programme is based. I conclude that this evidence for the preferability of intrasentential rhetorical marking is at best only mildly supportive, and at worst neutral. As I argue, the earlier heuristic proposed by the authors, where the number of propositions per sentence is limited, rather than maximised) seems preferable.

As for the aim that the markers used be as unambiguous as possible (an aim I examine throughout this chapter), I find it tells only part of the story: the pressures to make text coherent, unambiguous and stylistically well-expressed can conflict with each other, each motivating different choices of connectives. A good text planner needs to take all these pressures into account when choosing how to realise its rhetorical structure—not just the pressure for non-ambiguity.

Next, I look at the practical rules for text generation S&dS derive from these assumptions. I examine how much sense these rules make from a discourse-analytical and from a syntactic perspective. Some of their rules, I conclude, are well-motivated and linguistically sound—particularly those concerned with ELABORATION, which is usually handled with inadequate finesse in text generation. At other times, they make erroneous predictions and exaggerated claims—most clearly in the rule deciding where propositions should be embedded in the clause.

I believe the main reason for these shortcomings is that the authors invest too much power in RST as an analytical device (and are not alone in doing so.) RST is not a general theory of linguistics; it is not a general theory of syntax, nor of propositional semantics (see Chapter 2), nor even of discourse semantics—as there is so much textual coherence that does not fall under its intended analytical scope.

In the third and fourth parts of my investigation, I investigate issues involving characteristically hypotactically-signalled, and unsignalled rhetorical relations—both of which raise serious challenges to Scott & de Souza's assumptions in formulating a programme of rhetorically unambiguous text output. In particular, I investigate three matters: the sensitivity of rhetorical connectives to textual meaning and text structure (which is particularly high for hypotactic connectives, and which text planners ignore at their peril); whether there are any connectives that distinguish between the three relations expressible by *but* (CONTRAST, ANTITHESIS and CONCESSION) (the answer is yes, but the connectives involved are stylistically marked, and cannot be used with any generality); and possible explanations for the fact that certain rhetorical relations tend not to be textually signalled at all (mainly that the relations in present are very strongly cued by their pragmatic context.)

My discussion may give the impression of being a systematic hatchet job; yet my intent has not been to use S&dS as straw(wo)men, nor to denigrate their research—which, I believe, in places far outstrips that of their contemporaries'. Rather, I have aimed to illustrate the danger of computational linguists taking too cavalier an attitude towards integrat-

ing into their field results from natural-language linguistics—whether discourse analysis, psycholinguistics, formal semantics, or syntax. I do not believe that either the pitfalls I have pointed out in S&dS, or the circumstances that have brought them about, are unique to those researchers. Indeed, I believe such problems in integrating found-language–linguistic results into text generation are endemic to the venture; and more than anything else, this chapter is intended to serve as a cautionary tale.

While that may have been the primary intent behind this chapter, though, several suggestions and proposals are made in it, relevant either to text generation or to various discourse-motivated aspects of theoretical linguistics. Several of these, I believe, deserve further research; and I am hopeful that such research will help bring about a more solid and robust bond between the two disciplines.

5. Conclusion.

There are two somewhat disparate threads motivating this work. They are bound together by the fact that both involve problems with text generation and RST—on the one hand, (structural) problems *within* RST, highlighted by the use of RST in text generation; one the other, problems *occasioned* by RST, when text generation researchers attempt to integrate RST analysis with the text planning task.

The first thread motivates Chapters 2 and 3. As a result of my investigations in those Chapters, I have contributed to three ontological issues within RST, in the hope that it will become a more rigorous theory of text structure, and in order to make the computational use of RST more methodical and less *ad hoc*. In particular:

• I attempted to provide a more cogent, less arbitrary motivation for making rhetorically relevant distinctions between relations, within RST. To that end, I formalised the *relational criterion,* used by Sanders *et al.* (1992). I attempted to formulate the criterion as a rigorous test, and I explored how the use of this criterion could affect our view of rhetorical structure, and of how it harmonises with pragmatics and formal semantics.

• I pointed out the discrepancy in what is considered an atomic discourse entity between text-linguistic, computational, and psycholinguistic views of rhetorical structure. I attempted to formulate these differing views in more concrete, propositional-semantic and syntactic terms; and I tentatively suggested ways in which they may be reconciled. I also explored how a non-traditional view of rhetorical structure, such as seems to be espoused by computational researchers, may introduce new challenges to rhetorical theory, in removing some of the simplifications in the linguistic data made by the traditional version of the theory.

• I argued for the need for a well-defined taxonomy of RST, on both theoretical-linguistic and computational grounds. I critiqued several attempts (by discourse linguists of various schools) to taxonomise rhetorical theory relations. I then proposed my own taxonomy of rhetorical relations, which I believe has the following advantages over previous such attempts: it is comprehensive; it is feature-based; and it has an explicit external motivation, being argued extensively in terms of symbolic logic paraphrases of Mann & Thompson's relation definitions. I also demonstrated the usefulness of my taxonomical approach by using it to account for the rhetorical ambiguity of several connectives.

The second thread in my thesis—the problems occasioned by the computational usage of RST—motivates Chapter 4. Its main intent is to illustrate the problems that can arise when the results of RST (and of similar linguistic research) are applied in an overly cavalier function to the task of text generation. I make this point by examining in detail the text planning heuristics formulated by Scott & de Souza—their strengths, their shortcomings, and where they are at odds with results from within linguistic theory.

In making this point, my discussion has led me to several investigations concerning the textual signalling of rhetorical relations. These investigations are founded in a variety of linguistic disciplines: discourse semantics; Relevance and presupposition theory; cohesion theory; psycholinguistics and reading comprehension research; syntax; stylistics; and (tentatively) grammaticalisation and functionalism. While none of these investigations have the explanatory scope of their counterparts in Chapters 2 and 3, I believe they are interesting springboards for further research into both theoretical and computational linguistics.

My thesis has subjected Mann & Thompson's Rhetorical Structure Theory to an exhaustive critique from many fronts. I believe there are areas in RST which are underspec-

ified, underdefined, or underexploited. In many ways, I attempt to show RST in its worst light. However, I emphatically do not reject RST as an analytical technique. While RST is not a general theory of discourse semantics (a misconception I believe is at the base of many of the problems I have discussed), within its particular analytical domain it has gained a deserved pre-eminence in the last seven years. Likewise, its application to text generation has given impressive results, as well as motivating in the field a long-overdue sensitivity to discourse structure.

It is precisely because of the proven explanatory and generative power of RST, that I consider it essential for the theory to be made more rigorous—both in its theoretical infrastructure, and as the methodology used by computational researchers. It is my hope that such rigour will prevent the discourse-linguistic and the computational versions of RST from separating any further, and, indeed, will bring them closer together. Both fields have only to gain from such a prospect.

Appendix A: RST Relation Classification on Informational/Presentational Basis

(Definitions cited from Mann & Thompson 1987)

Informational relations

ELABORATION
Constraints on N: none.
Constraints on S: none.
Constraints on the N+S combination: S presents additional detail about the situation or some element of subject matter which is presented in N or inferentially accessible in N in one or more of the ways listed below. In the list, if N presents the first member of any pair, then S includes the second:

1. set: member
2. abstract: instance
3. whole: part
4. process: step
5. object: attribute
6. generalization: specific

Effect: R recognizes the situation presented in S as providing additional detail for N. R identifies the element of subject matter for which detail is provided.

CIRCUMSTANCE
Constraints on N: none.
Constraints on S: presents a situation (not unrealized.)
Constraints on the N+S combination: S sets a framework in the subject matter within which R is intended to interpret the situation presented in N.
Effect: R recognizes that the situation presented in S provides the framework for interpreting N.

SOLUTIONHOOD
Constraints on N: none.
Constraints on S: presents a problem.[117]
Constraints on N+S combination: the situation presented in N is a solution to the problem stated in S.
Effect: R recognizes the situation presented in N as a solution to the problem stated in S.

CONDITION
Constraints on N: none.
Constraints on S: S presents a hypothetical, future, or otherwise unrealized situation (relative to the situational context of S.)
Constraints on the N+S combination: Realization of the situation presented in N depends on the realization of that presented in S.
Effect: R recognizes how the realization of the situation presented in N depends on the realization of that presented in S.

OTHERWISE
Constraints on N: presents an unrealized situation.
Constraints on S: presents an unrealized situation.
Constraints on the N+S combination: realization of the situation presented in N prevents the realization of the situation presented in S.

---

[117] "In the definition of the solutionhood relation, the terms problem and solution are broader than one might expect. The scope of problem includes:

1. questions
2. requests, including requests for information
3. some descriptions of desires, goals, intellectual issues, gaps in knowledge or other expressions of needs
4. conditions that carry negative values, either expressly or culturally, including calamities and frustrations." (M&T:51–52)

Effect: R recognizes the dependency relation of prevention between the realization of the situation presented in N and the realization of the situation presented in S.

INTERPRETATION
Constraints on N: none.
Constraints on S: none.
Constraints on the N+S combination: S relates the situation presented in N to a framework of ideas not involved in N itself and not concerned with W's positive regard.
Effect: R recognizes that S relates the situation presented in N to a framework of ideas not involves in the knowledge presented in N itself.

EVALUATION
Constraints on N: none.
Constraints on S: none.
Constraints on the N+S combination: S relates the situation presented in N to degree of W's positive regard toward the situation presented in N.
Effect: R recognizes that the situation presented in S assesses the situation presented in N and recognizes the value it assigns.

RESTATEMENT
Constraints on N: none.
Constraints on S: none.
Constraints on N+S combination: S restates N, where S and N are of comparable bulk.
Effect: R recognizes S as a restatement of N.

SUMMARY
Constraints on N: none.
Constraints on S: none.
Constraints on the N+S combination: S presents a restatement of the content of N, that is shorter in bulk.
Effect: R recognizes S as a shorter restatement of N.

SEQUENCE
Constraints on N: multi-nuclear.
Constraints on the combination of nuclei: A succession relationship between the situations is presented in the nuclei.
Effect: R recognizes the succession relationship between the nuclei.

CONTRAST
Constraints on N: multi-nuclear.
Constraints on the combination of nuclei: no more than two nuclei; the situations presented in these two nuclei are (a) comprehended as the same in many respects (b) comprehended as differing in a few respects and (c) compared with respect to one or more of these differences.
Effect: R recognizes the comparability and the difference(s) yielded by the comparison being made.

VOLITIONAL CAUSE
Constraints on N: presents a volitional action or else a situation that could have arisen from a volitional action.
Constraints on S: none.
Constraints on the N+S combination: S presents a situation that could have caused the agent of the volitional action in N to perform that action; without the presentation of S, R might not regard the action as motivated or know the particular motivation; N is more central to W's purposes in putting forth the N–S combination than S is.
Effect: R recognizes the situation presented in S as a cause for the volitional action presented in N.

VOLITIONAL RESULT
Constraints on N: none.
Constraints on S: presents a volitional action or a situation that could have arisen from a volitional action.
Constraints on the N+S combination: N presents a situation that could have caused the situation presented in S; the situation presented in N is more central to W's purposes than is that presented in S.

Effect: R recognizes that the situation presented in N could be a cause for the action or situation presented in S.

NON-VOLITIONAL CAUSE
Constraints on N: presents a situation that is not a volitional action.
Constraints on S: none.
Constraints on the N+S combination: S presents a situation that, by means other than motivating a volitional action, caused the situation presented in N; without the presentation of S, R might not know the particular cause of the situation; a presentation of N is more central than S to W's purposes in putting forth the N–S combination.
Effect: R recognizes the situation presented in S as a cause of the situation presented in N.

NON-VOLITIONAL RESULT
Constraints on N: none.
Constraints on S: presents a situation that is not a volitional action.
Constraints on the N+S combination: N presents a situation that caused the situation presented in S; presentation of N is more central to W's purposes in putting forth the N–S combination than is the presentation of S.
Effect: R recognizes that the situation presented in N could have caused the situation presented in S.

PURPOSE
Constraints on N: presents an activity.
Constraints on S: presents a situation that is unrealized.
Constraints on the N+S combination: S presents a situation to be realized through the activity in N.
Effect: R recognizes that the activity in N is initiated[118] in order to realise S.

Presentational relations

MOTIVATION
Constraints on N: presents an action in which R is the actor (including accepting an offer), unrealized with respect to the context of N.
Constraints on S: none.
Constraints on the N+S combination: Comprehending S increases R's desire to perform action presented in N.
Effect: R's desire to perform action presented in N is increased.

ANTITHESIS
Constraints on N: W has positive regard for the situation presented in N.
Constraints on S: none.
Constraints on the N+S combination: the situations presented in N and S are in contrast (cf. CONTRAST, i.e., are (a) comprehended as the same in many respects (b) comprehended as differing in a few respects and (c) are compared with respect to one or more of these differences); because of an incompatibility that arises from the contrast, one cannot have positive regard for both the situations presented in N and S; comprehending S and the incompatibility between the situations presented in N and S increases R's positive regard for the situation presented in N.
Effect: R's positive regard for N is increased (through R's comprehension of the incompatibility of the situations presented in N and S.)

---

[118]Initiating an action is typically a volitional action, but PURPOSE can have a non-volitional nucleus. M&T discuss the implications of such non-volitionality by giving the following example:

> *'Presumably, there is a competition among trees in certain forest environments to be come as tall as possible* ***so as to catch as much of the sun as possible for photosynthesis.'***
>
> However, in all such examples [of non-volitional PURPOSE] that we have found or imagined, some purpose seems implied. That is, there is in the subject matter some tendency toward particular classes of outcomes or states, and the span that expresses purpose identifies those outcomes or states for which the tendency supposedly exists.
>
> In the examples just cited, the purpose clause implies a teleological perspective on anatomical attributes. From this perspective, trees are as they are because they are embedded in a framework in which organisms tend toward photosynthesis maximization. (M&T:65)

BACKGROUND

Constraints on N: R won't comprehend N sufficiently before reading text of S.

Constraints on S: none.

Constraints on the N+S combination: S increases the ability of R to comprehend an element in N.

Effect: R's ability to comprehend N increases.

ENABLEMENT

Constraints on N: presents R action (including accepting an offer), unrealized with respect to the context of N.

Constraints on S: none.

Constraints on the N+S combination: R comprehending S increases R's potential ability to perform the action presented in N.

Effect: R's potential ability to perform the action presented in N increases.

CONCESSION

Constraints on N: W has positive regard for the situation presented in N.

Constraints on S: W is not claiming that the situation presented in S doesn't hold.

Constraints on the N+S combination: W acknowledges a potential or apparent incompatibility between the situations presented in N and S; W regards the situations presented in N and S as compatible; recognizing the compatibility between the situations presented in N and S increases R's positive regard for the situation presented in N.

Effect: R's positive regard for the situation presented in N is increased (through recognition of the compatibility of situations presented in N and S.)

JUSTIFY

Constraints on N: none.

Constraints on S: none.

Constraints on the N+S combination: R's comprehending S increases R's readiness to accept W's right to present N.

Effect: R's readiness to accept W's right to present N is increased.

EVIDENCE

Constraints on N: R might not believe N to a degree satisfactory to W.

Constraints on S: The reader believes S or will find it credible.

Constraints on the N+S combination: R's comprehending S increases R's belief of N.

Effect: R's belief of N is increased.

Appendix B: Rhetorical Taxonomies.[119]

B.1. Longacre's Rhetorical Taxonomy

Basic

Conjoining

Coupling (JOINT): *He's short and he's fat.*

Contrast (CONTRAST): *I don't like hamburgers, but my wife does.*

Comparison (*COMPARISON*): *John loves Mary more than he loves Susan.*

Alternation (*DISJUNCTION*): *Either he did it or he didn't.*

Temporal

Overlap (CIRCUMSTANCE): *As he prayed, he walked alone.*

Succession (SEQUENCE): *They played tennis for an hour, then swam for another hour.*

Implication

Conditionality (CONDITION): *If large doses of vitamin C are harmful, I'll stop taking it.*

Causation (CAUSE cluster, CIRCUMSTANCE): *You didn't go because you were afraid.*

Contrafactuality (CONDITION): *If he hadn't gone, I wouldn't have complained.*

Warning (OTHERWISE): *We shouldn't let our torches go out, otherwise we'll never find our way home.*

Elaborative [120]

Paraphrase

Information-preserving

Equivalence (RESTATEMENT): *He capitulated immediately; he surrendered on the spot.*

Negated Antonym (CONTRAST): *It's not black, it's white*.

Information-increasing

Generic-specific (ELABORATION): *He was executed yesterday, he was shot by the firing squad.*

Amplification (ELABORATION): *He went away, I saw him go away.*

Information-decreasing

Specific-generic (RESTATEMENT): *They dug up Assyrian ruins, they did some excavation.*

Contraction [121](SUMMARY): *I won't go to see him, I just won't go.*

Summary (SUMMARY): *John works at the sawmill; Jim at the repair shop; and Al at the print shop—that's what they're all doing.*

Illustration

Simile (*COMPARISON*): *Go tell Herod who is like a fox.*

Exemplification (ELABORATION): *He has had an innovating career as seen in his introduction of the Mariachi Mass into the Cathedral.*

---

119 The RST relations which given taxa are closest to are given in brackets and SMALL CAPS. RST relations M&T considered, but did not formulate in their inventory, are *ITALICISED*.

120"[Elaborative relations] are considered to be essentially embellishments, i.e., rhetorical devices. […] The elaborative devices are a further extension of the elementary statement calculus." (Longacre 1983:80) Compare: "An elaborating relation is used when there is a relation of 'being' between two or more units; this is the very general relation that obtains between an attribute and a value, between a set and its members, or between a generalization and its specific instances. Elaborating relations are distinct from other rhetorical relations in that they do not necessarily hold between propositions per se, but may relate terms in the propositions, e.g. one term in the proposition may be related to another as type to subtype. Like the 'rhetorical act' [Presentational] relations […] elaborating relations also tend to be scale-insensitive; they may occur at any level in a rhetorical structure." (Matthiessen & Thompson 1988:298)

121"But while in specific–generic paraphrase the loss of information in the second base is due to the use of a more generic lexical term than in the first base, in contraction paraphrase, certain lexical items (often noun phrases) which are found in first base are not found in the second base at all." (Longacre 1983:121)

Deixis

Introduction (BACKGROUND?): *There was a young man named Amkidit, he lived on the mountain. And…*

Identification (BACKGROUND?): *Kimboy went back and got a hammer and that was what they used.*

Attribution

Speech (*QUOTE*): *"I'm fine," said John, "But how are you?"*

Awareness (*QUOTE*)**:** *I know that he's coming.*

Frustration [122] (CONTRAST, ANTITHESIS, CONCESSION): *She's fat but she's not sloppy.*

RST relations not accounted for in Longacre's scheme: SOLUTIONHOOD, INTERPRETATION, EVALUATION, MOTIVATION, ENABLEMENT, JUSTIFY.

B.2. Halliday & Hasan's Rhetorical Taxonomy.

Additive

External/Internal: Additive, simple

Additive (JOINT): *and*

Negative (JOINT): *nor*

Alternative (ALTERNATIVE): *or*

Internal

Complex

Emphatic

Additive (JOINT): *furthermore*

Alternative (ALTERNATIVE): *alternatively*

De-emphatic: *Afterthought* (JOINT): *by the way*

Apposition

Expository (RESTATEMENT): *in other words*

Exemplificatory (EVIDENCE): *for instance*

Comparison

Similar (*COMPARISON*): *likewise*

Dissimilar (ANTITHESIS): *on the other hand*

Adversative

External/Internal: Adversative 'proper' (ANTITHESIS, CONTRAST)

Simple: yet, though, only

Containing and*: but*

Emphatic: however

External: Contrastive (CONTRAST)

Simple: but, and

Emphatic: however

Internal

Contrastive: *Avowal* (ANTITHESIS): *in fact*

Correction (ANTITHESIS)

Of Meaning: instead

Of Wording: I mean

Dismissal (CIRCUMSTANCE)

Closed: in either case

Open-ended: however it is

Causal

External/Internal

Causal, General (CAUSE Cluster)

Simple: so

Emphatic: consequently

Causal, Specific

---

[122] "Many of the structures described in previous sections have frustrated counterparts. […] Basic to this notional structure is an implication that there is a P which normally implies a Q, but that rather than Q the opposite positive–negative value occurs." (Longacre 1983:134)

Reason (CAUSE): *for this reason*
Result (RESULT): *as a result*
Purpose (PURPOSE): *with this in mind*

Conditional

Simple (CONDITION): *then*
Emphatic (CONDITION): *in that case*
Generalized (CONDITION): *under the circumstances*
Reverse Polarity (OTHERWISE): *otherwise*

Internal

Reversed Causal: *Simple* (CAUSE): *because*
Causal, Specific (EVIDENCE)
Reason: it follows
Result: arising out of this
Purpose: to this end
Respective (—)
Direct: in this respect
Reversed Polarity: aside from this

Temporal

External/Internal: Conclusive: *Simple* (SEQUENCE): *finally*

External

Temporal, Simple (SEQUENCE)
Sequential: then, next
Simultaneous: just then
Preceding: before that
Temporal, Complex (CIRCUMSTANCE)
Immediate: at once
Interrupted: soon
Repetitive: next time
Specific: next day
Durative: meanwhile
Terminal: until then
Punctiliniar: at this moment

Internal

Internal Temporal (PRESENTATIONAL SEQUENCE)
Sequential: secondly
Conclusive: finally, in conclusion
'Here and Now' (—)
Past: up to now
Present: at this point
Future: from now on
Summary
Summarizing (SUMMARY): *in short*
Resumptive (—): *to resume*

RST relations not accounted for in Halliday & Hasan's scheme: SOLUTIONHOOD, INTERPRETATION, EVALUATION, MOTIVATION, ENABLEMENT, JUSTIFY, ELABORATION, BACKGROUND, RESTATEMENT, CONCESSION.

B.3. The Tilburg Rhetorical Taxonomy.

Causal

Semantic

Positive: *Cause-consequence* (VOLITIONAL CAUSE, NON-VOLITIONAL CAUSE, VOLITIONAL RESULT, NON-VOLITIONAL RESULT): *so that, because/in consequence of (the fact that)*

Negative: *Contrastive cause-consequence* (CONCESSION): *although/despite the fact that, but; hoewel, maar*[123]

[123]Dutch 'prototypical markers' of the relations.

Pragmatic

Positive: *Argument-claim* (EVIDENCE): *so, therefore, since, for, because*
Instrument-goal (PURPOSE): *(in order) to, so as to, to that end*
Condition-consequence (CONDITION): *provided that*

Negative: *Contrastive argument-claim* (CONCESSION): *although, but; al, maar*[1]

Additive

Semantic

Positive: *List* (JOINT, SEQUENCE): *and, also*

Negative: *Exception* (CONTRAST): *but*
Opposition (CONTRAST): *but, however, by contrast*

Pragmatic

Positive: *Enumeration* (*PRESENTATIONAL SEQUENCE*): *moreover*

Negative: *Concession* (CONCESSION, CONTRAST): *but, yet; maar*[1]

RST relations not accounted for in the Tilburg scheme: SOLUTIONHOOD, INTERPRETATION, EVALUATION, MOTIVATION, ENABLEMENT, JUSTIFY, ELABORATION, ANTITHESIS, CIRCUMSTANCE, OTHERWISE, RESTATEMENT, SUMMARY, BACKGROUND.

Appendix C: Logical Symbols Index.

Denotational semantics:
⟦A⟧: the denotation of A.
Pragmatics:
+>: implicates.
S_p: agent *p* says utterance (locution) S.

The following symbols are taken from Gensler (1990):
Modal Logic:
◊A: A is logically[124] possible (true in some possible world.)
Imperative (deontic) logic:
A: Do A! Make A true!
A(u): u, make A true! u should bring about A.
OA: It ought to be that A; Act A is obligatory (required, mandatory, a duty)
RA: It would be all right that A; Act A is right (permissible, OK)
Belief (doxastic) logic:
u:A: u accepts that A; u believes that A.
u:A: u accepts that A should happen; u wills A or does A.
u:A ∧ A: u wills A and A happens; u does A.

[124]In the context of this work, this operator is weakened from *logically* possible to *physically* possible. It does not mean 'morally possible, allowed', which is RA.

Appendix D. Maier & Hovy's taxonomy.

I only obtained a copy of Maier & Hovy (1993) after I had worked out the taxonomy detailed in Section 3.6. I believe it is of interest to compare the two.

Maier & Hovy (1993) follows on from earlier work by Hovy in 1990, where he collected 350 relations from the work of 25 researchers, and merged them into a taxonomy. Maier & Hovy reorganise Hovy's 1990 results; in particular, at the highest level of the taxonomy, relations are divided according to Halliday's three metafunctions of language: ideational, interpersonal, and textual. As often discussed in my thesis, the ideational/interpersonal distinction seems to motivate the presentational/informational distinction. In contrast to how I view textual meaning, however, Maier & Hovy do not consider textual meaning to correspond to Elaborative relations. Instead, they seem to take a similar view of 'coherence in the world' as did Hobbs (1985), and end up classing as 'textual' relations like CONJUNCTION, DISJUNCTION, and PRESENTATIONAL SEQUENCE.

Maier & Hovy's taxonomy is given in Fig. D.1.

The main methodological fault I find with Maier & Hovy's taxonomy is that it is not a feature-based analysis. A fully-fledged, classical componential analysis of relations is probably unduly restrictive, as we have seen with Sanders *et al.* (1992), and it is doubtful it would be of much help in classifying Elaborative relations. However, a limited feature-based analysis, such as that I have attempted, does help capture generalisations *between* categories, which a simple tree-structure like Maier & Hovy's misses.

Some of the better known clusters of relations turn up in Maier & Hovy's scheme. The CAUSE cluster is there, as are the presentational members of the Causation cluster. Many of the groups M&T identified are also here: INTERPRETATION and EVALUATION, RESTATEMENT and SUMMARY; temporal relations are collapsed into CIRCUMSTANCE à la M&T, etc. In fact, although it contains several relations not in M&T (COMPARISON, ANALOGY, EXCEPTION, QUALIFICATION), Maier & Hovy's scheme includes all relations present in M&T (unlike any other of the taxonomies I considered.)

In some details, Maier & Hovy's classifications differ considerably from mine. They reject a distinct Elaborative class in their taxonomy (which was present in Hovy's earlier scheme.) But the function of ELABORATION and CIRCUMSTANCE, and for that matter BACKGROUND, is quite different to that of 'coherence in the world' relations like CAUSE, CONDITION, and SEQUENCE. In fact, it seems to me peculiar that relations like ELABORATION would not be considered to bear textual meaning, but relations like DISJUNCTION would. In addition, Maier & Hovy persist with M&T's classification of BACKGROUND as a Presentational relation; as I have argued, BACKGROUND is out of place in that class.

Interestingly, Maier & Hovy class SOLUTIONHOOD, INTERPRETATION and EVALUATION as Presentational relations. As I have argued, whether SOLUTIONHOOD is Presentational or not depends on its definition. If it is Presentational, it would occupy the same niche in my taxonomy as ENABLEMENT. In itself, that is not a problem: JUSTIFY and EVIDENCE also occupy the same niche, being distinguished by Sweetser's Epistemic/Illocutionary parameter.

But if SOLUTIONHOOD does inhabit the same niche as ENABLEMENT, it would be worth asking whether the distinction between them is relational or not. As I discussed above, the distinction between the two is whether there is an explicit request involved. This is not *per se* a property of the relation between the two related text spans; but neither is it a property of the referential meaning of the text, since it involves an illocution.

Maier & Hovy also deviate from M&T in classing EVALUATION and INTERPRETATION as Presentational. But neither has the explicit intent to persuade the reader of anything, or to do any action. Rather, they interpret and organise the text, and thereby help in building up the mental model of the discourse. Since these relations are bound to the text *qua* text so explicitly, I have classed them as Elaborative relations, conveying textual rather than interpersonal meaning. But this depends on point of view; Systemic-Functional Linguistics does consider affective meaning part of Interpersonal meaning, and if these relations are considered primarily affective, rather than text-organisational, then a stronger case could be made that they are Presentational.

Another way in which Maier & Hovy are at odds with my classification is their classifying OTHERWISE with CONTRAST and COMPARISON, in a different class than CONDITION. This seems particularly counterintuitive, since OTHERWISE is associated so strongly with CONDITION; indeed, as I have argued, it is associated with CONDITION more closely than it is with any adversative relation.

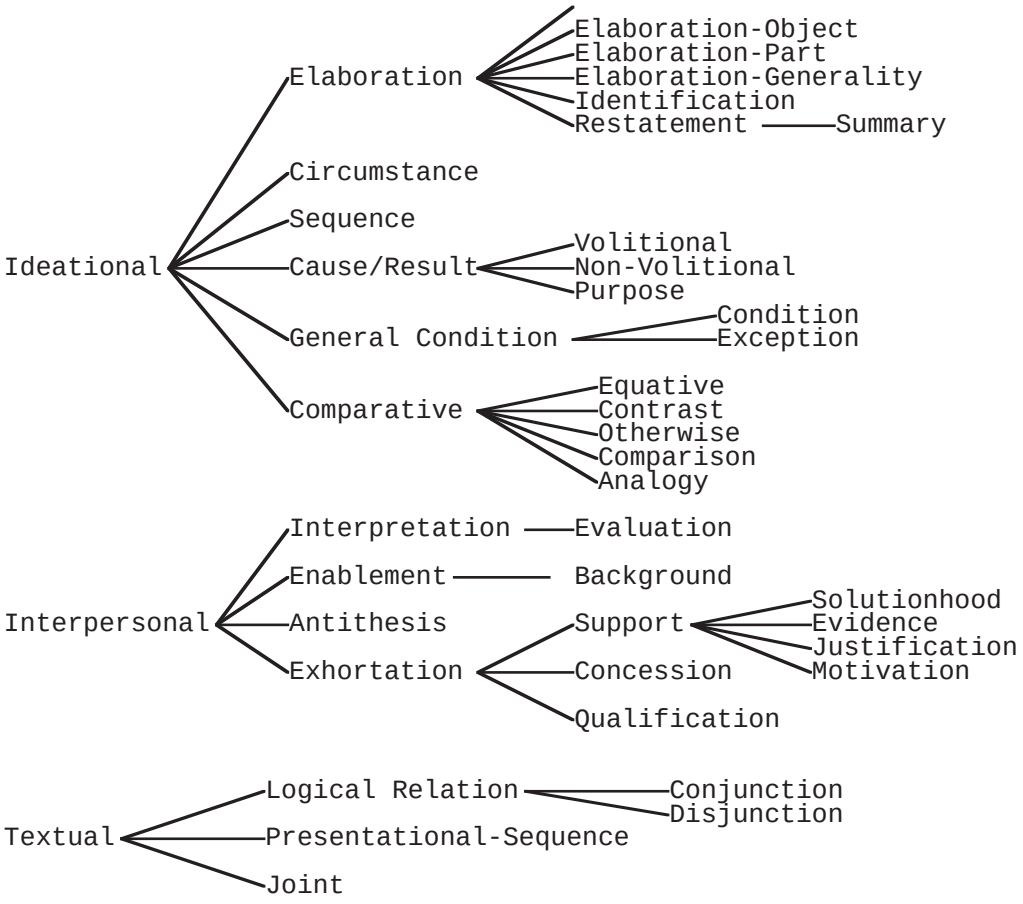

Fig. D.1. Maier & Hovy's taxonomy.

Appendix E. Proof that CONDITION is Modal.

At least prototypically, CONDITION makes no claim that its nucleus is true (though this is not given as a defining constraint in RST.) Longacre makes the presence of such a claim the differentiating factor between CAUSE and CONDITION:

> Causation, as we have said, involves not simply an implication, but a given. That is, there is not only an antecedent consequent relation, but the antecedent is factual or is at least assumed to be so for the sake of the argument. (Longacre 1983:106)

But, if ⟦S⟧∧(⟦S⟧⇒⟦N⟧) holds (which is how Longacre describes Causation), then ⟦N⟧ holds. So CAUSE must imply that its nucleus is true, and this is how I have been logically translating CAUSE: CAUSE*(S,N)*≡(⟦N⟧∧ (⟦S⟧⇒⟦N⟧).)

Now by contrast, RST sets on the satellite of CONDITION the constraint that it presents "a hypothetical, future, or otherwise unrealized situation". Is the antecedent of CONDITION always logically possible? Logicians have long illustrated paradoxes by using antecedents that are not; e.g. *If 1+1=3, then I'm a monkey's uncle*.

But McCawley (1993) rightly points out that most such examples are contrived, failing relevance logic: it is anomalous in natural language for the antecedent and consequent of an implication not to have any real-world relevance to each other (being a monkey's uncle has nothing to do with arithmetic.)

There are few instances in natural language usage where a logical impossibility like *1+1=3* (as opposed to a contingent, physical impossibility—like me handing in my thesis by the due date, when I am typing these words two days after that date) is relevant in a CONDITION. Now, if something is not *logically* impossible, there is a possible (hypothetical, but logically consistent) world in which that something is true. But this is the same as to say that it is logically possible: ◊*Something*.

So we can conclude that there always exists some possible world (maybe even our own), in the antecedent of any CONDITION expression is possible: ◊⟦S⟧. But by modal logic, ◊⟦S⟧∧⟦S⟧⇒⟦N⟧ imply ◊⟦N⟧. So CONDITION makes the claim required (◊⟦N⟧) for it to be considered Modal.[125]

[125]CONDITION involves logical possibility, rather than the stronger condition of physical possibility invoked by the other modal relations. I still believe it valid to use this weaker of the two notions of possibility to group these relations together.

Bibliography.

Anderson, R. C. & Davison, A. 1988. Conceptual and Empirical Bases of Readability Formulas. In Davison & Green.

Appelt, D. E. 1985. *Planning English Sentences*. Cambridge: Cambridge University Press.

Baker, E. L., Atwood, N. K. & Duffy, T. M. 1988. Cognitive Approaches to Assessing the Readability of Text. In Davison & Green.

Bishop, J. 1993. *Rhetorical Structure Theory and Advertisements*. Honours Thesis (B.A.) University of Melbourne.

Blakemore, D. 1989. Denial and Contrast: a Relevance Theoretic Analysis of *But*. *Linguistics and Philosophy* **12**: 15–37.

Blakemore, D. 1993. Understanding Utterances: An Introduction to Pragmatics. Oxford: Blackwell.

Chafe, W. 1988. Linking intonation units in spoken English. In Haiman & Thompson.

Dale, R.; Hovy, E.; Rösner, D. & Stock, O. (eds.) 1992. *Aspects of Automated Natural Language Generation*. Lecture Notes in Artificial Intelligence 587. Berlin: Springer.

Dale, R.; Mellish, C. & Zock, M. (eds.) 1990. *Current Research in Natural Language Generation*. London: Academic Press.

Davison, A. & Green, G. M. (eds.) 1988. *Linguistic Complexity and Text Comprehension: Readability Issues Reconsidered*. Hillsdale (New Jersey): Lawrence Erlbaum Associates.

de Beaugrande, R.–A. & Dressler, W. 1981. *Introduction to Text Linguistics*. London: Longman.

de Souza, C. S., Scott, D. R. & Nunes, M. de G. V. 1989. Enhancing Text Quality in a Question-Answering System. In Martins & Morgado.

Dowty, D. R., Karttunen, L. & Zwicky, A. M. 1985. *Natural language parsing: Psychological, computational, and theoretical perspectives*. Studies in Natural Language Processing 1. Cambridge: Cambridge University Press.

Eggins, S. 1994. *An Introduction to Systemic Functional Linguistics.* London: Pinter.

Fletcher, C. R. 1981. Short-term memory processes in text comprehension. *Journal of Verbal Learning and Verbal Behavior* **20**:546–574.

Flood, J. (ed.) 1984. *Understanding Reading Comprehension*. Delaware: International Reading Association.

Flores d'Arcais, G. B. 1978. The Perception of Complex Sentences. In Levelt & Flores d'Arcais.

Fox, B. 1987. *Discourse Structure and Anaphora*. Cambridge: Cambridge University Press.

Frazier, L. 1985. Syntactic Complexity. In Dowty *et al.*

Gensler, H. J. 1990. *Symbolic Logic: Classical and Advanced Systems*. Englewood Cliffs (New Jersey): Prentice-Hall.

Graesser, A. C.; Hoffman, N. L. & Clark, L. F. 1980. Structural Components of Reading Time. *Journal of Verbal Learning and Verbal Behavior* **19**:135–151.

Grosz, B. J. & Sidner, C. L. 1986. Attention, Intentions and the Structure of Discourse. *Computational Linguistics* **12**:175–204.

Grimes, J. E. 1975. *The Thread of Discourse*. Janua Linguarum Series Minor 207. The Hague: Mouton.

Haiman, J. & Thompson, S. A. (eds.) 1988. *Clause Combining in Grammar and Discourse*. Amsterdam: John Benjamins.

Halliday, M. & Hasan, R. 1976. *Cohesion in English*. London: Longman.

Hasan, R. 1984. Coherence and cohesive harmony. In Flood.

Hobbs, J. R. 1985. On the Coherence and Structure of Discourse. Report No. 85–37. Stanford (California): Center for the Study of Language and Information, Stanford University.

Hoey, M. 1991. *Patterns of Lexis in Text*. Describing English Language Series. Oxford: Oxford University Press.

Horacek, H. & Zock, M. (eds.) 1993. *New Concepts in Natural Language Preparation.* London: Pinter.

Hovy, E. H. 1990. Unresolved Issues in Paragraph Planning. In Dale *et al.*

Hovy, E. H. 1991. Approaches to the Planning of Coherent Text. In Paris *et al.*

Hovy, E. H. 1993. From interclausal relations to discourse structure—a long way behind, a long way ahead. In Horacek & Zock.

Jarvella, R. J. 1970. Effects of syntax on running memory span for connected discourse. *Psychonomic Science* **19**:235–236.

Jarvella, R. J. 1971. Syntactic Processing of Connected Speech. *Journal of Verbal Learning and Verbal Behavior* **10**:409–416.

Kintsch, W. & van Dijk, T. A. 1978. Toward a model of text comprehension and production. *Psychological Review*. **85**:363–394.

Knott, A. & Dale, R. 1993. Using Linguistic Phenomena to Motivate a Set of Rhetorical Relations. Technical Report HCRC/RP–39, Human Communication Research Centre, University of Edinburgh.

Krifka-Dobeš, Z. & Novak, H.–J. 1993. From constituent planning to text planning. In Horacek & Zock.

Levelt, W. J. M. & Flores d'Arcais, G. B. 1978. *Studies in the Perception of Language*. Chichester: John Wiley & Sons.

Longacre, R. E. 1983. *The Grammar of Discourse*. New York: Plenum.

Maier, E. & Hovy, E. 1993. Organising discourse structure relations using metafunctions. In Horacek & Zock.

Mann, W. C. 1987. *Text Generation: The Problem of Text Structure*. New York: Springer.

Mann, W. C., & Thompson, S. A. 1986. Relational Propositions in Discourse. *Discourse Processes* **9**: 57–90.

Mann, W. C. & Thompson, S. A. 1987. *Rhetorical Structure Theory: a Theory of Text Organisation*. ISI Reprint Series ISI/RS–87–190. Marina del Rey (California): Information Sciences Institute.

Martins, J. P. & Morgado, E. M. (eds.) 1989. *EPIA 89: 4th Portugese Conference on Artificial Intelligence Proceedings*. Lecture Notes in Artificial Intelligence 390. Berlin: Springer.

Matthiessen, C. & Thompson, S. A. 1988. The Structure of Discourse and 'Subordination'. In Haiman & Thompson.

McCawley, J. 1993. *Everything that linguists have always wanted to know about logic but were ashamed to ask*. Chicago: University of Chicago Press.

McKeown, K. R. 1985. *Text Generation: Using Discourse Strategies and Focus Constraints to Generate Natural Language Text*. Studies in Natural Language Processing. Cambridge: Cambridge University Press.

Moore, J. D. & Paris, C. L. 1993. Planning Text for Advisory Dialogues: Capturing Intentional and Rhetorical Information. *Computational Linguistics* **19**: 651–694.

Moore, J. D. & Pollack, M. E. 1992. A Problem for RST: The Need for Multi-Level Discourse Analysis. *Computational Linguistics* **18**: 537–544.

OED (Simpson, J. A. & Weiner, E. S. (eds.)) 1989. *The Oxford English Dictionary*. Oxford: Clarendon Press.

Paris, C. L.; Swartout, W. R. & Mann, W. C. (eds.) 1991. *Natural Language Generation in Artificial Intelligence and Computational Linguistics*. Boston: Kluwer Academic Publishers.

Quirk, R.; Greenbaum, S.; Leech, G.; & Svartvik, J. 1972. *A Grammar of Contemporary English*. London: Longman.

Reichman, R. 1985. *Getting Computers to Talk Like You and Me: Discourse Context, Focus, and Semantics (An ATN Model.)* Cambridge (Mass.): MIT Press.

Rösner, D. & Stede, M. 1992. Customising RST for the Automatic Production of Technical Manuals. In Dale *et al.*

Rutherford, W. E. 1970. Some observations concerning subordinate clauses in English. *Language* **46**: 97–115.

Sanders, T. J. M.; Spooren, W. P. M. & Noordman, L. G. M. 1992. Toward a Taxonomy of Coherence Relations. *Discourse Processes* **15**: 1–35.

Scott, D. R. & de Souza, C. S. 1990. Getting the Message Across in RST-based Text Generation. In Dale *et al.*

Shadbolt, N. R. 1989 [1986]. Planning and Discourse. In Taylor *et al.*

Simonin, N. 1988. An Approach for Creating Structured Text. In Zock & Sabah.

Stede, M. 1992. Linearizing RST Trees in the Bilingual Generation of Instructional Text. Technical Report 92020, FAW Ulm (Germany.)

Sweetser, E. 1990. *From Etymology to Pragmatics: Metaphorical and cultural aspects of semantic structure*. Cambridge Studies In Linguistics 54. Cambridge: Cambridge University Press.

Taylor, M. M.; Neel, F. & Bouwhuis, D. G. (eds.) 1989 [1986]. *The Structure of Multimodal Dialogue*. Human Factors in Information Technology 4. Amsterdam: North Holland.

Thompson, S. A. & Mulac, A. 1991. A Quantitative Perspective on the Grammaticization of Epistemic Parentheticals in English. In Traugott & Heine.

Traugott, E. C. & Heine, B. (eds.) 1991. *Approaches to Grammaticalization*. Amsterdam: Benjamins.

van Dijk, T. A. & Kintsch, W. 1983. *Strategies of Discourse Comprehension*. New York: Academic Press.

van Dijk, T. A. 1977. *Text and Context: Explorations in the Semantics and Pragmatics of Discourse*. Longman Linguistics Library 21. London: Longman.

Vander Linden, K. 1993. Ph.D. Dissertation. University of Colorado Tech Report CU–CS–654–93.

Vander Linden, K.; Cumming, S. & Martin, J. 1992. Using System Networks to Build Rhetorical Structures. In Dale *et al.*

Widdowson, H. G. 1978. *Teaching Language as Communication.* Oxford: Oxford University Press.

Zock, M. & Sabah, G. (eds.) 1988. *Advances in Natural Language Generation: An Interdisciplinary Perspective*, Vol. 1. London: Pinter.