

# Parameters for Rhetorical Structure Theory Ontology<sup>1</sup>

Nick Nicholas, University of Melbourne

Rhetorical Structure Theory (RST), a framework for analysing discourse structure above the clausal level, has been used extensively by both text analysts and researchers in text generation for the past decade. The proliferation of relations posited by various researchers using RST has highlighted an ontological problem with the theory: there have been no well-defined criteria determining which types of relations should belong within the framework, and what degree of granularity should be used to distinguish between these relations. Sanders, Spooren & Noordman (1992) have proposed the *relational criterion* as a criterion to perform this task. I examine the implications of applying their criterion to temporal relations, as a case study, and the implications the criterion may have for RST as a whole.

## 1. What is Rhetorical Structure Theory?

This paper deals with the ontology of Rhetorical Structure Theory (RST) (Mann & Thompson (1987)—henceforth M&T). RST is a framework developed to account for text structure above the clause level; it does so by positing hierarchical relations between spans of text. It is used by linguists not only as a tool for analysing the structure of natural language text, but also as a planning aid in text generation—in particular, to help the computer decide on means to linguistically realise these intratextual relations.

RST uses a number of predefined rhetorical relations. These relations are defined functionally, in terms of what their intended effect on the reader is. Examples of such relations are JUSTIFY, ELABORATION, PURPOSE, ANTITHESIS, and CONDITION. Most relations are asymmetrical, with a *nucleus* span differentiated from *satellite* spans. While syntactic subordination is only one way of realising a rhetorical relation, it is easiest to think of the nucleus as equivalent to a matrix clause, and the satellite as equivalent to a qualifying adjunct. The full definition of a rhetorical relation consists of constraints on the text spans related; constraints on the combined span; and the Effect—a description of the relation's expected effect on a reader.

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. 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)

According to their Effect, rhetorical relations are subdivided by M&T 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 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 represented by a truth-conditional semantics. Examples of such relations are ELABORATION, CIRCUMSTANCE, PURPOSE, CONDITION, and SUMMARY.

---

<sup>1</sup>This paper (which is work in progress) is based on the first half of chapter 2 of my Master’s thesis (Nicholas 1994). My thanks to Dr Lesley Stirling for her continuing support and comments, and to John Bowden, Drs Robert Dale and Brian Paltridge, Andrew Malone and Ghodrattollah Kamyab for their valuable insights.

*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 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.




Fig. 1. RST analysis of ‘Music Day’ text.

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.

Note that CAUSE is *not* a ‘semantic’ relation because of the way it is inferred to be present in a text. A CAUSE relation can be implicated just as well as it can be entailed. CAUSE is, rather, semantic (and this is, in a sense, the main point of this paper) because a formal representation of CAUSE requires only the tools of formal semantics. The representation of CAUSE can address facts in the ‘outside’ world, without necessitating a model of the speaker and the hearer. No formal representation of MOTIVATION, on the other hand, can fail to incorporate such a speaker-hearer model: the illocution and perlocutions of the relation need to be modelled carefully, and demand a formal Speech-Act theory. In *that* respect, CAUSE is Informational/Semantic (being perlocutionally trivial), and MOTIVATION is Presentational/Pragmatic (being perlocutionally complex.)

## 2. The Problem

### 2.1. The Problem in General

A major problem with RST has become apparent, particularly with recent computational work in RST (e.g. Scott & de Souza 1990, Rösner & Stede 1992, Maier & Hovy 1993), which has arbitrarily expanded the inventory of relations used. This problem which has caused much discussion within computational linguistics circles.

The problem is, what kinds of relations should be included in the inventory of RST? In particular, what should the balance be between Informational relations, which are of a formal-semantic nature, and Presentational relations, which are primarily of a pragmatic nature? To answer this question, we need to establish what type of theory RST is to be. Is a full propositional semantics a component of RST? For that matter, is a full Speech Act Theory a component of RST?

Since the dichotomy of Presentational (pragmatic) and Informational (semantic) relations is a well-established feature of RST, present since its inception, little attention has been paid to this underlying issue—despite the resulting problems (some of which I discuss below.) However, there are strong hints in Mann & Thompson’s earlier writing on rhetorical theory, that the relations should behave in a way associated with pragmatics, rather than semantics.

1. Relational propositions<sup>2</sup> are relatively ‘basic’, in the sense that there is a tendency for many other sorts of inferences to be derived from them, but they tend not to be derived from other sorts of inferences.
2. Relational propositions arise in a text independently of any *specific* signals of their existence.
3. Relational propositions are involved in communicative (or illocutionary) ‘acts’, in the sense of ‘speech acts’. (Mann & Thompson 1986:68)

But even if we accept, on the basis of this, that rhetorical relations are primarily pragmatic in nature, this still does not address what level of granularity a rhetorical theory should use in discriminating between relations. Nor does it answer how RST should treat undeniably semantic/informational relations (such as CAUSE), which nonetheless few rhetorical analysts would wish to discard.

This question becomes particularly relevant when RST is extended to encompass intrasentential relations—as frequently occurs in text generation work (Scott & de Souza 1990, Rösner & Stede 1992, Krifka-Dobeš & Novak 1993, Vander Linden 1993). This is because relations between text spans *within* a sentence are likelier to be not pragmatic, but informational/semantic. That is to say, they are less likely to be defeasible, to invoke situational context, or to bear a nontrivial perlocutionary effect—all hallmarks of pragmatic relations—and more likely to be explicitly signalled textually, lexically conventionalised and entailed, and to bear only the trivial perlocutionary effect of informing the reader about some fact in the real world. This, of course, is because intrasententially linked clauses are linked by explicit connectives like *because* and *to*, while intersentential rhetorical linking affords the option of no explicit connective, or of a much vaguer connective (such as *and*—which can introduce anything from JOINT and SEQUENCE clauses to CAUSES, MOTIVATIONS, and BACKGROUNDS.)

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 do not. To give a concrete example, I consider Rösner & Stede’s (1992) claim that the inventory of rhetorical relations should include an UNTIL relation. I evaluate various approaches towards circumscribing the ontology of RST—and defining what relations belong and don’t belong—by looking at how they handle the UNTIL relation. To begin with, I summarise what I consider two extremes in the approaches formulated to date. I then examine Sanders *et al.*’s (1992) proposal of the *relational criterion* to address this issue. I consider how they apply the criterion; I point out some problems with their approach, and give an alternative approach to the relational criterion and the place of temporal relations in RST, drawing on Lascarides, Asher & Oberlander’s (1992) model of inferring temporal relations in a text.

## 2.2. A Specific Case of the Problem: UNTIL

Consider: what should a rhetorical theory do with a text like *Heat the pot until it starts frothing*? There is obviously a coherence relation<sup>3</sup> between two clauses in this text: the

---

<sup>2</sup>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)

<sup>3</sup>*Coherence relations* refers to all the relations between entities inside a text which contribute to a text being coherent; rhetorical relations are necessarily a subset of these relations. For example, coherence relations can apply between constituents in a noun phrase; the bottom level for rhetorical analysis appears to be daughter-nodes of S, although computational linguists have extended the scope of their analyses to clausal adjuncts (Nicholas 1994). On the other hand, a relation like BOTH SENTENCES ARE CONCERNED WITH THE VIRTUES OF SCOTTISH PIZZA may pass as a coherence relation, but does not contribute significantly to our understanding of the rhetorical structure of a text (a point Knott & Dale (1993a) make forcefully.) From a rhetorical point of view, it may be preferable to view the relation as an instance of ELABORATION or JOINT.

interpretation of the anaphor *it* depends on the earlier *the pot*; and there is a clear syntactic relation of hypotaxis between two clauses—characteristic of most rhetorical relations. But how should the coherence relation be specified in a rhetorical theory—that is, using the tools of which module of linguistics, and to what level of refinement?

As Knott & Dale (1993a) point out, there are two logical extremes for a rhetorical theory. On the one extreme, there are only two rhetorical relations: Presentational and Informational. On the other, each instance of a rhetorical link in any text becomes its own taxon, and there are infinitely many relations. Rhetorical analysts would regard both alternatives as untenable; but there is a wide space between the two extremes, along which rhetorical analysts can select their alignment. Depending on this alignment, the relation above can be analysed as a CIRCUMSTANCE, as in M&T; a temporal relation, as in Hobbs (1985); or, specified even further, as a punctiliniar temporal relation, as in Halliday & Hasan (1976), or an UNTIL relation, as in Rösner & Stede (1992).<sup>4</sup>

The difficulty many rhetorical analysts find with more inclusive approaches to what relations belong under the compass of RST is that RST itself becomes murky and unfalsifiable. When RST becomes all-inclusive, its predictive power is compromised. (See Knott & Dale (1993a, 1993b), Nicholas (1994) for arguments.) In this particular case, to put it informally, text analysts don't believe a rhetorical theory should be in the temporal logic business. That is, the distinction between UNTIL and, say, a relation like DURING is made (at length) in temporal logic, a formal semantics. If such distinctions are admitted into RST wholesale, there will 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, 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. This seems imprudent: 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. (Indeed, with no principled ontology, and a strictly empirical basis for deriving its repertoire of relations, RST fails to meet the epistemological criteria of a theory.)

So does a rhetorical theory need to include a temporal logic component? M&T seem to think it does not, since they collapse all temporal relations in their scheme into CIRCUMSTANCE—a relation not even principally temporal. Most researchers involved in rhetorical analysis would probably intuitively agree with their decision. But until recently, there has been no attempt to make an explicit formulation of this intuition (which could be used to generate a judgement on why PURPOSE is admissible in RST, but UNTIL is not.) Nor have the wider implications of this intuition for the theory been adequately explored.

### 3. Approaches to the Problem

As already noted, the issue of which relations belong in RST has drawn much attention in the past few years, and several researchers have formulated approaches to addressing it. I summarise below two well-known approaches which represent opposite poles in treating this issue, and illustrate how each would treat the case of UNTIL.

---

<sup>4</sup>For M&T's definition of CIRCUMSTANCE, and Rösner & Stede's definition of UNTIL, see Appendix. Note that, in the 'fine print' of Rösner & Stede's paper, UNTIL is associated with Imperative nuclei (this follows from the instructional types of texts they are generating)—although this is not made explicit in their definition. This deontic version of UNTIL is somewhat more complex semantically than the non-deontic version; I will address it at the end of this paper.

### 3.1. Hovy: ‘Profligate’

Eduard Hovy has been a long-time advocate of RST in text generation, and has been working with RST since his PhD thesis in 1988. In Hovy (1991) (which differs little from his subsequent work, Maier & Hovy (1993)), he formulates an explicit defense for what he terms the ‘profligate’ approach to RST ontology (as opposed to the ‘parsimonious’ approach taken in Grosz & Sidner (1986)). According to the ‘profligate’ approach, the rhetorical inventory should include tens of relations, and indeed should remain (as originally stated in M&T) open-ended.

To ensure the inventory does not become unmanageable, Hovy allows for *incremental* refinement of divisions between relations, in a hierarchical framework. Thus, at the top level of his hierarchy (following Halliday (1985)), he distinguishes between ELABORATIONS (relations between entities and their properties or components), ENHANCEMENTS (circumstances of place, time, manner etc.), and EXTENSIONS (ontologically more distant relations—causes, contrasts etc.) At the first level, ENHANCEMENT can be broken up into CIRCUMSTANCE, BACKGROUND, and SOLUTIONHOOD. CIRCUMSTANCE in turn can be broken up into LOCATION, TIME, MANNER, and so on. Hovy supports his scheme by surveying 25 researchers on what relations they feel belong in a rhetorical theory, and placing these relations into his schema.

The objection raised to relations like UNTIL is that they are too formal-semantic and not rhetorical enough to belong in a rhetorical theory. Hovy’s response is,

[a]s a result of this study and previous work, I believe [...] that no *purely* rhetorical or *purely* semantic relations exist. Some relations, typically those higher in the hierarchy, certainly fulfill a more rhetorical function (i.e., provide more information about the argumentational structure of the discourse) than a semantic function. But not even the topmost relations are entirely devoid of semantic meaning. (Hovy 1991:134)

As to the fact that this approach does introduce the full complexity of formal semantics into a rhetorical theory, Hovy willingly acknowledges this—and does not consider it a serious problem,

as long as the terms are well-behaved and subject to a pattern of organization which makes them manageable [...]. Each discourse relation simply inherits from its ancestors all necessary processing information, such as cue words and realization constraints, adding its unique peculiarities, to be used for inferring its discursal role (in parsing) or for planning out a discourse (in generation). Increasing differentiation of relations, continued until the very finest nuances of meaning are separately represented, need be pursued only to the extent required for any given application. (Hovy 1991:134)

Thus, under Hovy’s schema, UNTIL is quite licit as a rhetorical relation; presumably, it would be counted as a special case of the TIME relation, itself a special case of CIRCUMSTANCE.

The notion of hierarchical semantic refinement of rhetorical relations is appealing; indeed, it is implicit in the very wording of the question I asked above, “what level of granularity should a rhetorical theory use in discriminating between relations?” I do not dispute Hovy’s determination that there is a gradient of coherence relations—the more relations are subdivided, the more semantic the features used to divide the relations up become—and the less relevant to textual organisation (i.e. the less rhetorical) they become.

However, Hovy’s approach sidesteps the problem I have alluded to, rather than resolving it. In particular, while no one would deny that UNTIL is a semantic special case of a temporal CIRCUMSTANCE, Hovy gives no account as to why none of the researchers he surveyed do consider UNTIL as a distinct relation, nor as to why certain relations are much more popular with researchers than others (e.g. CAUSE/RESULT (14 researchers) as against QUALIFICATION (one researcher).) Hovy does not address what the semantic/rhetorical threshold should be, because he feels there is no such threshold. But this does not accord with the intuitions of most rhetorical analysts; nor are their concerns really resolved by his approach. So while the rhetorical/semantic spectrum is indeed something of a continuum, researchers are justified in believing a possibly arbitrary threshold needs to

be imposed on it, below which any further subdivision of relation is rhetorically irrelevant, and pertinent only to ‘hard’ semantics.

The other major problem with Hovy’s approach is that he fails to address the Moore & Paris (1993) problem (discussed immediately below) of the many-to-many relation between information and intentionalities in text. The fact that this relation is many-to-many means that the problem these researchers highlight—that some RST relations have to do with intentionality, while others have to do with information, and neither can be mapped onto the other—cannot be resolved solely by a hierarchical refinement model—which can only represent one-to-many mappings.

### 3.2. Moore et al.: *Stratification*

As already noted, RST in its original formulation incorporates relations of both an informational and a presentational or intentional nature—relations, in other words, to do with both the state of the world and speaker intentions. While M&T admit the possibility of ambiguity in RST analyses of text, they do not consider either this, or the conflation of two disciplines in the theory, as a problem. Within computational linguistics circles at least, Moore & Pollack (1992) have firmly established that this conflation is in fact a problem. Indeed, the current crisis RST is undergoing in the computational linguistics community was precipitated by their article.

Moore & Pollack (1992) establish that texts such as *George Bush supports big business; he’s sure to veto House Bill 1711* can be plausibly interpreted as involving either an informational relation (CAUSE), or a presentational relation (EVIDENCE). They support the contention made by other discourse analysts (such as Grosz & Sidner (1986)) and systemic linguists, that both the informational and presentational aspects of coherence (in Hallidayan terms, the ideational and the interpersonal) are co-present in any text. RST claims to cover both domains, but constrains any one textual analysis to select only one relation as holding between any two text spans. Therefore, they believe, RST has serious problems both computationally and descriptively; and these problems cannot be readily fixed.

A further problem is highlighted in Moore & Paris (1993). While for presentational relations the intention underlying the relation is isomorphic to the relation itself, informational relations are in a many-to-many relationship to intentions. There is no way to determine uniquely (isolated from context) the intention behind uttering the CIRCUMSTANCE *The screwdriver’s in the top drawer of the toolbox*; nor is there a unique rhetorical realisation of the intention ‘Enable the hearer to find the screwdriver.’ (Thus instead of a CIRCUMSTANCE, one could use a CONTRAST: *Not the yellow one.*)

As already argued, this problem in particular cannot be resolved by Hovy’s approach. Indeed, the fact that neither the informational nor the presentational can be reduced to the other has led some analysts—notably Korelsky & Kitredge (1994)—to argue for a stratification of RST. In their particular model, informational relations are treated as *realisations* of intentions, and the authors propose conditioned rules should be found for mapping intentions to informational relations, inasmuch as there is no isomorphism. (Such work is also adumbrated in Moser & Moore (1993).)

Under such an approach, UNTIL is squarely delimited as an Informational relation, and has no interaction with Presentational relations, other than as a realisation of various intentions. That is to say, a relation within a text can be either UNTIL, DURING, or CAUSE, say—from an informational point of view—and it can be either MOTIVATION, ENABLEMENT, or JUSTIFY—from a presentational point of view—but it cannot be either UNTIL or MOTIVATION, or either CAUSE or JUSTIFY: presentational and informational relations are not allowed into the same rhetorical inventory. Whenever UNTIL is encountered in a text, there is always an underlying presentational relation present, and realised by UNTIL. But by the imposition of only a *mapping* relation between informational and presentational relations, instead of allowing a paradigmatic relation as in M&T’s RST, the two become completely distinct types of entities. In contrast with Hovy’s approach, where the rhetorical/semantic threshold is totally absent, here it becomes a chasm.

The stratification approach has various wide-ranging consequences, which I will not explore fully here. It would certainly result in a theory of text totally unlike M&T’s RST,

although possibly not incompatible with Speech Act theory. A particular feature I find of concern is that, in Korelsky & Kittredge's model, overt textual realities are related only to the informational level of text; intentional relations are abstract, and are not affected by e.g. linear textual order or discourse connectives. It is well known that there are discourse connectives with an overtly Presentational rather than Informational function: in English, *as a result* (CAUSE) vs. *it follows that* (EVIDENCE) (Knott & Dale 1993a); in French, *parce que* (CAUSE) vs. *puisque* (EVIDENCE); in Dutch, *al* (PRAGMATIC CONCESSION—cf. *he isn't sick although he hasn't shown up for work today*) vs. *hoewel* (SEMANTIC CONCESSION—cf. *he isn't sick although he had gone skinny-dipping last night*) (Sanders *et al.* 1992). So a stratified RST would give an impoverished account of such connectives: if intentionality is not reflected directly in text, then how is the difference between *al* and *hoewel* to be explained by a stratified RST?

My concern in this paper is to find a way of formalising the intuition that CAUSE belongs in a rhetorical theory, while UNTIL does not. A stratification approach would take us too far afield: it would necessarily insulate both CAUSE and UNTIL from Presentational relations like JUSTIFY and MOTIVATION. An all-inclusive approach like Hovy's, on the other extreme, fails to address the feeling that some types of relation are legitimately considered rhetorical and contribute to the organisation of text, whereas others are merely semantic. While the Moore & Pollack problem is real enough, there is still a place for an intuitive, rule-of-thumb analysis of text (conflating the ideational and the interpersonal), such as is provided by RST. But if its results are to have any significance, RST still needs to be formalised more than it has been to date. I believe the *relational criterion*, as a middle road between the two extreme approaches discussed, provides just such a means of formalisation.

#### 4. The Relational Criterion

In their attempt to establish parameters according to which coherence relations can be classified, Sanders *et al.* (1992) (see also Sanders & Spooren (1990), Sanders (1993)) set up what they call 'the relational criterion'. Their intent is to generate a parsimonious and cognitively plausible taxonomy of coherence relations—which they attempt to support with psycholinguistic data. This taxonomy uses binary features to distinguish between relations; the parameters they consider for this role need to satisfy the relational criterion to be accepted into the taxonomy:

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)

To paraphrase: say we have two coherence relations, which differ only by the presence or absence of a semantic or pragmatic feature. Is this feature rhetorically relevant? That is, should the two relations be considered two distinct rhetorical relations, or two instances of the one rhetorical relation, the difference between which is not relevant to a rhetorical theory? The answer is: the feature is rhetorically relevant iff it passes the relational criterion. It passes the relational criterion, in turn, if the 'subject matter' of the feature is not the meaning of the nucleus or satellite text span taken in isolation, but the meaning *added* to the text by positing that coherence relation. (As Sanders *et al.* term it, the 'informational surplus'.)

If the 'subject matter' of the feature is the semantics of the nucleus or satellite *in isolation*, then the feature involves semantic distinctions which are quite independent of the coherence relation. Whether the particular relation is in fact present or absent in the text has no bearing on whether the feature holds or not, since the feature follows from the clausal semantics, and not from any coherence relation. So if the feature does not involve

the relation between the nucleus and satellite, it can hardly be considered relational. But if the feature is not relational, it has no place in a taxonomy of rhetorical relations.

In the version of the relational criterion I use, I take the wording of the relational criterion to its logical conclusion (this step is certainly implicit in Sanders *et al.*'s treatment of temporals): if the feature considered concerns meaning following *deductively* from the meaning of the nucleus and satellite spans taken in isolation, then it is not rhetorically relevant. Such a feature involves: the isolated meaning of the text spans; the juxtaposition of the text spans (and hence the knowledge that they are, at least, in a coherence relation); and the rules of formal-semantic inference. The only 'informational surplus' here is the fact that there exists a coherence relation; if the particular relation involved—say, AFTER—follows directly by deduction, then its choice has not *added* anything to the meaning of the text. Indeed, the question of choice of the particular relation never arises; it is necessarily true by entailment. The relation does not shape the meaning of the text, but merely follows from the semantics of the individual clauses in juxtaposition.

By contrast, a feature *will* pass the relational criterion iff it involves a difference in the relation between text segments *qua implicated* (defeasible) meaning. 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—and not by the semantics of the individual text spans—then presupposing another rhetorical relation will eliminate this feature from the text meaning; so the added information is defeasible.

(Note that this defeasibility can only be the case where the 'informational surplus' is *not* manifest or overtly signalled in the text, e.g. by connectives. The semantics of connectives such as *since* is, of course, entailed, and cannot be defeated and substituted by, say, an ANTHESIS reading.)

For example, the fact that the text

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

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 a generalisation. ('You can't teach an old dog new tricks. For example, I'm not going to start learning Dutch.')

But if CAUSE is substitutable by EVIDENCE in the absence of any overt textual relational marking, then the distinction between CAUSE and EVIDENCE is defeasible. So it does not follow by entailment from the semantics of the constituent text spans; rather, the fact that there is a CAUSE rather than an EVIDENCE relation is 'informational surplus'. Therefore, the difference between CAUSE and EVIDENCE passes the relational criterion, and is rhetorically relevant.

Now if, on the other hand, we attempt to posit a different coherence relation in a context where that relation is entailed by the text span semantics, then we will get nonsense: entailments are not defeasible. Take for example the sentence *I flicked the switch and the light came on*. We know that this sentence involves a CAUSE relation. M&T would further classify this as a VOLITIONAL CAUSE. I contend (Nicholas 1994) that the volitionality of this relation is entailed by the semantics; the evidence for this is that the presupposition of volitionality cannot be defeated: ??*I flicked the switch on and, even though I didn't mean for it to happen, the light came on.*<sup>5</sup>

While I do not agree with the particular rhetorical taxonomy Sanders *et al.* have formulated, it is clear that the relational criterion offers a bold and sound solution to our prob-

---

<sup>5</sup>Admittedly, a context can be manufactured in which this makes sense (I've rigged the fuse box so that the light wasn't supposed to come on.) This possibility is also the case for UNTIL relations; and I address this problem later in the paper.



lem: any relations that (being defeasible) lie within the analytic domain of pragmatics—whether Speech Act Theory proper, or conversational implicature theory—are considered rhetorically relevant. Any relations that (being entailed) lie outside pragmatics, and are directly subject to a truth-conditional semantic analysis, are automatically excluded from rhetorical consideration. Note that, unlike the stratification approach, the relational criterion does *not* exclude all informational relations from the rhetorical inventory. As illustrated above, the distinction between CAUSE—an Informational relation—and EVIDENCE—a Presentational relation—passes the criterion.

Furthermore, the informational relations rejected are rejected not because of an *a priori* divide imposed between the pragmatic and the semantic, but because they do not behave in a way consistent with the theory. The implicated, defeasible nature of features passing the criterion is more consistent with Mann & Thompson's (1986) description of relational propositions, than are the truth-conditionally semantic features that fail it. Indeed, the requirement made by the relational criterion, that the relation should follow from the combination of the text spans rather than the individual text spans, is explicit in Mann & Thompson's formulation:

The type of text-conveyed information that we are interested in [...] is the *relational proposition*, which arises (in the hearer's tacit awareness) [...] from two parts of a text but is not independently derived from either of these parts. The phenomenon is, in other words, a *combinational one*, defined on two parts of a text. (Mann & Thompson 1986:58)

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. This is precisely what the relational criterion achieves.

One may well ask: if Mann & Thompson made the requirement for the combinational nature of rhetorical relations explicit, why do we need the relational criterion at all? There are three reasons. First, Mann & Thompson's statement is less rigorous, and much looser, than the relational criterion proper: relations entailed by clausal juxtaposition are admissible (although Mann & Thompson do make it explicit in the same paper that defeasibility is a characteristic feature of relational propositions.) Second, M&T themselves do not follow the relational criterion, including in their RST taxonomy such distinctions as CAUSE and RESULT, and VOLITIONAL CAUSE and NON-VOLITIONAL CAUSE. Third, the criterion has been followed even less by those who have continued M&T's work, leading to the current crisis in RST; therefore, a rigorous explicit formulation of the relational criterion is necessary.

#### 4.1. *The Relational Criterion against Temporals*

Sanders *et al.* (1992) argue that temporal relations (including UNTIL—and, indeed, SEQUENCE) do not belong in a rhetorical theory. That is to say, while temporal relations are very much coherence relations in text, a rhetorical theory should not go so far as to make any distinctions between them and simple additive relations, like JOINT:<sup>6</sup>

[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

---

<sup>6</sup>Sanders *et al.* (1992) treat the distinctions between rhetorical relations as feature distinctions (as does Nicholas (1994)). Under such an analysis, the most 'unmarked' relation—the relation for which the component features all have their unmarked value—is JOINT: simple conjunction. (It is no accident that M&T do not even consider JOINT a distinct relation, but a rhetorical schema—in effect, merely a textual juxtaposition.) Because of this unmarkedness, JOINT is liable to end up a grab-bag for coherence relations not otherwise distinguished by the rhetorical features in the taxonomy—as occurs here with temporals. Note that Sanders *et al.* do not consider temporal relations as CIRCUMSTANCES, since they have chosen to exclude Elaborative relations from their taxonomy (unlike Nicholas (1994).)

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)

So consider an UNTIL relation, like *They turned the radiator cap counter-clockwise until it stopped*. We can take it as given that the nucleus *they turned the radiator cap* and the satellite *until it stopped* 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.

But 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 (or for that matter, CAUSE) could still be implicated between the two; it is only world-knowledge that makes us select UNTIL.

In fact, let us use an example where there is no explicit textual signal of the temporal relation to begin with. In ‘Max fell. John pushed him’, the temporal relation is one of PRECEDENCE (Max fell after John pushed him, by Cause and Effect.) In the following text, however, the same text spans are related in a temporal relation of SEQUENCE:

John and Max came to the cliff’s edge. John applied a sharp blow to the back of Max’s neck.  
*Max fell. John pushed him.* Max rolled over the edge of the cliff.

Now, since the text spans are identical in both cases, there is no question of PRECEDENCE and SEQUENCE being entailed from different individual text spans. For the difference between PRECEDENCE and SEQUENCE to pass the relational criterion, it must concern the ‘informational surplus’ added to the text by positing the new relation. Sanders *et al.* would argue that the temporal qualities of the events described follow from the individual segments. Clearly, however, they are part of the ‘informational surplus’ brought to bear into the text by changing the posited relation: whether the first event precedes or follows the second depends entirely on pragmatic factors—context and real-world knowledge. So it seems the distinction between PRECEDENCE and SEQUENCE passes the relational criterion after all.

This is not the answer we expected—it is certainly not the answer Sanders *et al.* came up with. 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. But the temporal information is not necessarily linguistically realised. Therefore, it is as subject to implicature and defeasibility as causation is; it too can form part of the ‘informational surplus’.

Then again, bringing up PRECEDENCE as a candidate rhetorical relation is something of a leap itself—although there is nothing in the relational criterion proscribing this relation. To resolve this matter, I will appeal to an explicit model of discursual inferencing—something Sanders *et al.* have not done. This will help better motivate why temporal relations should be considered as entities distinct from rhetorical relations. It will also show precisely how temporal information does form part of the ‘informational surplus’.

## 5. Temporal distinct from Rhetorical: Another Approach

In their work, Lascarides, Oberlander and Asher have provided an explicit theory on how temporal relations can be implicated in a text, and how this implication interacts with both real-world and rhetorical information. While it does not itself provide a solution to the problem I consider, it does give a clear illustration of precisely how temporal relations fit into the pragmatic/semantic divide.

Lascarides & Asher (1991) have developed DICE (Discourse and Commonsense Entailment), a nonmonotonic [defeasible] logic framework for discourse representation, which incorporates hierarchical structuring of discourse, and rhetorical relations.<sup>7</sup> This framework can be used to infer both the rhetorical relations holding in a text, and the temporal relations between the events described in the text. This is done by using a default logic (e.g.  $\langle \tau, \alpha, \beta \rangle \text{NARRATION}(\alpha, \beta)$ :<sup>8</sup> if text span  $\beta$  is to be linked to text span  $\alpha$  in text  $\tau$ , assume *by default* that they are in a NARRATION relation); certain axioms (e.g.  $\text{NARRATION}(\alpha, \beta) \rightarrow e_{\alpha} \leq e_{\beta}$ : if  $\alpha$  is linked to  $\beta$  as a NARRATION, the event in  $\alpha$  must temporally precede that in  $\beta$ ); and certain rules of deduction (like the defeasible Modus Ponens,  $\phi \triangleright \psi, \phi \vdash \psi$ : e.g. Birds normally fly, Tweety is a bird, so (by default) Tweety flies).

Lascarides *et al.* (1992) use the DICE framework to show how context interacts with the default assumption that time moves forwards in text (i.e. that the relation between text flow and time is by default iconic.) What is significant for my purposes is the epistemological status given time in their framework: the rhetorical relation comes first, and the temporal relation is *secondarily* inferred. This should become clear from the following rules in their scheme:

$\langle \tau, \alpha, \beta \rangle \text{NARRATION}(\alpha, \beta)$ : if text span  $\beta$  is to be linked to text span  $\alpha$  in text  $\tau$ , then assume by default that they are in a NARRATION relation.

$\text{NARRATION}(\alpha, \beta) \rightarrow e_{\alpha} \leq e_{\beta}$ : if  $\alpha$  is linked to  $\beta$  as a NARRATION, then the event in  $\alpha$  must temporally precede that in  $\beta$ .

$\langle \tau, \alpha, \beta \rangle \wedge \text{cause}(e_{\beta}, e_{\alpha}) \triangleright \text{EXPLANATION}(\alpha, \beta)$ : if text span  $\beta$  is to be linked to text span  $\alpha$  in text  $\tau$ , and the event in  $\beta$  caused that in  $\alpha$ , then assume by default that they are in an EXPLANATION relation.

$\text{EXPLANATION}(\alpha, \beta) \rightarrow \neg e_{\alpha} \leq e_{\beta}$ : if  $\alpha$  is linked to  $\beta$  as an EXPLANATION, then the event in  $\alpha$  cannot temporally precede that in  $\beta$ .

$\langle \tau, \alpha, \beta \rangle \wedge \text{cause}(e_{\alpha}, e_{\beta}) \triangleright \text{RESULT}(\alpha, \beta)$ : if text span  $\beta$  is to be linked to text span  $\alpha$  in text  $\tau$ , and the event in  $\alpha$  caused that in  $\beta$ , then assume by default that they are in an RESULT relation.

$\text{RESULT}(\alpha, \beta) \rightarrow e_{\alpha} \leq e_{\beta}$ : if  $\alpha$  is linked to  $\beta$  as a RESULT, then the event in  $\alpha$  must temporally precede that in  $\beta$ .

$\langle \tau, \alpha, \beta \rangle \wedge \text{state}(e_{\beta}) \triangleright \text{overlap}(e_{\alpha}, e_{\beta})$ : if the text span  $\beta$  to be linked to  $\alpha$  describes a state, then assume the events in  $\alpha$  and  $\beta$  overlap.

$\langle \tau, \alpha, \beta \rangle \wedge \text{overlap}(e_{\alpha}, e_{\beta}) \triangleright \text{BACKGROUND}(\alpha, \beta)$ : if text span  $\beta$  is to be linked to text span  $\alpha$  in text  $\tau$ , and the event in  $\beta$  overlaps with that in  $\alpha$ , then assume by default that they are in a BACKGROUND relation.

$\text{BACKGROUND}(\alpha, \beta) \rightarrow \text{overlap}(e_{\alpha}, e_{\beta})$ : if  $\alpha$  is linked to  $\beta$  as an EXPLANATION, then the events in  $\alpha$  and  $\beta$  must temporally overlap.

So a clear notion emerges in the DICE work that rhetorical relations (both presentational and informational) and real-world relations (both temporal and causal) are distinct entities. Since the rules mapping between them are defeasible and context-dependent, they are also

<sup>7</sup>Note, however, that the discourse relation inventory they use is Hobbs' (1985), which differs from M&T's. Thus, Lascarides & Asher (1991) use NARRATION (where RST has SEQUENCE), EXPLANATION (where RST has CAUSE), BACKGROUND, RESULT, and ELABORATION. Since they consider the inferencing of temporal relations, Lascarides *et al.* (1992: 2) do not feel they need to extend this rhetorical inventory any further, although they do admit "full coverage of text would require a larger set of relations, akin to that in [M&T]."

<sup>8</sup>For explanation of symbols used, see end of paper.

clearly irreducible. This is even the case where rhetorical and real-world relations may bear the same name, as in  $result(e_\alpha, e_\beta) \equiv cause(e_\beta, e_\alpha)$  and  $RESULT(\alpha, \beta)$ :  $RESULT$  (the rhetorical relation) always implies  $result$  (the semantic relation), but  $result$  does *not* always imply  $RESULT$ .

So how is the ‘John and Max’ text mentioned in §4.1 handled by DICE—how is the ‘informational surplus’ of the temporal properties of the clauses introduced into the discourse? Lascarides *et al.*’s derivation goes something like this:

- ‘John and Max came to the cliff’s edge’ is linked by NARRATION to ‘John applied a sharp blow to the back of Max’s neck’, which in turn is linked by NARRATION to ‘Max fell.’
- The only ‘open’ clause in the current text, therefore (that is, the only clause to which a new clause can be attached in a discourse structure), is ‘Max fell.’
- We wish to attach the clause ‘John pushed him’.
- Real-world knowledge gives us the Push Causal Law:  $\langle \tau, \alpha, \beta \rangle \wedge fall(m, e_\alpha) \wedge push(j, m, e_\beta) > cause(e_\alpha, e_\beta)$ . In English: if clause  $\alpha$  precedes clause  $\beta$  in text, and in clause  $\alpha$  Max falls while in clause  $\beta$  John pushes Max, then assume by default that the pushing in  $\beta$  caused the falling in  $\alpha$ .
- Furthermore, our knowledge of discourse gives us the Maintain Causal Trajectory Law (MCT): for some rhetorical relation  $R$ ,  $\langle \tau, \beta, \gamma \rangle \wedge R(\alpha, \beta) \wedge cause(e_\alpha, e_\beta) > \neg cause(e_\gamma, e_\beta)$ . In English: assume by default that the text does not jump around between cause and effect and back again: that is, rule out by default a *Cause-Result-Cause* text.<sup>9</sup>
- We had earlier decided that John applying a sharp blow to Max’s neck caused Max to fall.
- Therefore, according to MCT, John pushing Max presumably did *not* cause Max to fall, since that would give a *Cause-Result-Cause* text.
- But according to the Push Causal Rule (PCL), John pushing Max presumably *did* cause Max to fall.
- Therefore we have a conundrum: we cannot decide either way. (Lascarides *et al.* call this a ‘Nixon Diamond’, after their illustration: Quakers are pacifists, republicans are not pacifists, Nixon is both a quaker and a republican, so we cannot tell whether Nixon is a pacifist or not.)
- But Lascarides *et al.* argue that, whenever such a conundrum arises, discourse interpretation should select whichever relation least disrupts the existing pattern of relations in the discourse. This is their notion of ‘Inertia’, which they associate with Hobbs’ (1985) notion of Genre. Since NARRATION is the only rhetorical relation to have been used in the text so far, we select it to link ‘John pushed him’ into the text.
- It is axiomatic in DICE that, if two clauses are in a NARRATION relation, their events are in a relation of temporal sequence. Ergo, Max fell before John pushed him.

The derivation given illustrates the following:

- The decision on real-world properties of text, such as temporal precedence, follows from both real-world knowledge (PCL; Cause Implies Temporal Sequence rule) and knowledge about discourse conventions (MCT; Discourse Genre Inertia; definitions of rhetorical relations).
- Temporal knowledge, in particular, is not necessarily contained in individual text spans; nor is it necessarily explicitly signalled textually. As shown, it can often be a defeasible consequent inferred from pragmatic and rhetorical information. Therefore, it can legitimately be considered part of the ‘informational surplus’.
- However, temporal relations are clearly considered distinct from rhetorical relations, and neither type of relation can be reduced to the other.
- There is no real difference in the inferencing process between the way causality and temporality are handled: ‘information surplus’ is at work for both types of real-world information.

<sup>9</sup>This is, of course, a defeasible rule; the easiest way to defeat it is by explicit textual signal, as in ‘Max switched off the light. The room went pitch dark, *since he had* drawn the blinds *too*.’ Cf. ‘Max switched off the light. The room went pitch dark. He drew the blinds’, which would *not* normally be interpreted as *Cause-Result-Cause*. (The time of the third event is not specified as preceding the second—as it would be by a ‘had drawn’ pluperfect—but in fact, nor is it specified as occurring *after* it; the simple past in ‘drew’ is simply non-specific about this. The second text thus has no overt textual signal.)

## 6. A Reformulation of the Relational Criterion

The relational criterion as stated allows a rhetorical feature to pass if it arises pragmatically, as an informational surplus to the individual text spans—rather than being deducible from the individual text spans. As seen, this does not work as required: temporal precedence and sequence do arise as informational surplus; and any definition of the ‘meaning’ of individual text spans which allows for the subsequent deduction of temporal relations would be at the least anomalous.

There is, however, a straightforward way of modifying the criterion so as to allow us to maintain that temporal distinctions do not belong in a rhetorical theory. This is by asserting that, while the presence or not of the feature considered may arise in a text pragmatically—either through real-world information or knowledge of discourse conventions—it will still be rejected by the relational criterion if it is *reducible* to a feature specific to the individual text spans. While time information about events is as subject to implicature as is causal information (as demonstrated above), an UNTIL relation can still subsequently be reduced to segment-specific information, as follows:

$$\begin{aligned} \text{UNTIL}(e_1, e_2) \equiv & \exists t_1 \exists t_2: \text{time}(e_1) = [t_1, t_2] \\ & \text{which is temporal information specific to the nucleus event} \\ & \wedge \exists t_3: \text{time}(e_2) = [t_2, t_3] \\ & \text{which is temporal information specific to the satellite event} \end{aligned}$$

(In English: Event 1 stops just before  $t_2$ , and Event 2 starts at  $t_2$ .)<sup>10</sup>

Note that this equivalence is a material *iff*, and is not defeasible. That is, the UNTIL relation follows purely deductively from the text-span-specific time information, and vice versa. But if  $a \equiv b$ , then  $a$  is reducible to  $b$  and vice versa. It may well be the case that we know this information about the time contours of the two events because of pragmatics and discourse theory: it may have arisen, as ‘informational surplus’, in a way similar to the DICE-based inference illustrated above. Wherever the information came from, however, it will still be considered non-rhetorical if it can then be *reduced* to segment-specific information.

No such comparable reduction can apply to CAUSE: there is no conceivable property specific to either  $e_1$  or  $e_2$ , such that it can follow that  $\text{cause}(e_1, e_2)$ . So CAUSE must pass the relational criterion, since it is inherently relational, and an informational surplus to the information of the particular segments. On the other hand, even if UNTIL can be implicated as an informational surplus to what is apparent in the individual text spans, it can then be *reduced* to information specific to each text span; so it fails the criterion.

The following analogy may help to illustrate: consider that sentence  $\alpha$  is a box which contains red and yellow LEGO blocks, and sentence  $\beta$  is a box which contains blue and green LEGO blocks. Red and blue blocks are temporal information; yellow and green blocks are non-temporal, real-world information contained by the sentences. Deduction is

<sup>10</sup>The deontic version of UNTIL, which Rösner & Stede (1992) actually invoke, is somewhat more complex, but it still consists in its entirety of text-span-specific information. It can be formulated as follows:

$$\begin{aligned} \text{DEONTIC-UNTIL}(e_1, e_2) \equiv & \exists t_2 \exists t_3: \text{time}(e_2) = [t_2, t_3] \\ & \text{which is temporal information specific to the nucleus event} \\ & \wedge t_2 > t_{\text{now}} \\ & \text{which is temporal information specific to the satellite event} \\ & \wedge (\neg e_2 \wedge t_{\text{now}} < t < t_2) \leftrightarrow e_1 \end{aligned}$$

In English: Event 2 happens from  $t_2$  on, where  $t_2$  is in the future; and while  $t_2$  has not *yet* arrived, do event 1. (The condition  $t < t_2$  is to prevent necessitating event 1 again after  $e_2$  has finished.) The final requirement can be reduced to text-span-specific information as follows: For  $t = t_{\text{now}}$ , we know that  $\neg e_2$ , and therefore,  $e_1$ . Furthermore, orders become effective as of the time of speaking; therefore,  $e_1$  starts at  $t_{\text{now}}$ . Similarly, for  $t = t_2$ , we know that  $e_2$ , and therefore  $\neg e_1$ . So  $\text{time}(e_1) = (t_{\text{now}}, t_2)$ . Thus,  $(\exists t_2 \exists t_3: \text{time}(e_2) = [t_2, t_3] \wedge t_2 > t_{\text{now}} \wedge (\neg e_2 \wedge t < t_2) \leftrightarrow e_1) \rightarrow \text{time}(e_1) = (t_{\text{now}}, t_2)$ .

So DEONTIC-UNTIL( $e_1, e_2$ ) can be reduced to a combination of UNTIL( $e_1, e_2$ ) and (as the deontic component) CONDITION( $e_2, e_1$ ). The temporal component of DEONTIC-UNTIL, making it distinct from the CONDITION relation, is reducible to UNTIL, which we have established can be entailed by text-span-specific information. Therefore, the difference between DEONTIC-UNTIL and CONDITION is not rhetorically relevant.

the process of building bigger LEGO blocks from smaller blocks. The relational criterion states that, if a given big block (a rhetorical relation feature) can be built out of the blocks inside the boxes, without using glue or blocks of other colours (pragmatic inferencing), then the feature is not rhetorically relevant. For example, a tower of alternating yellow and green blocks can be built out of the contents of the boxes: the tower corresponds to a semantically entailed coherence relation feature, and is not rhetorically relevant. A tower of alternating yellow and mauve blocks cannot be built out of the box contents, since there are no mauve blocks in the boxes (the mauve is supplied by the discourse context, say.) Therefore this tower corresponds to the ‘informational surplus’ introduced into the text by a rhetorically relevant feature.

Now, according to Sanders *et al.*’s formulation of the criterion, temporal relations are towers of red and blue blocks. But Sanders *et al.* assume that the boxes always contain red and blue blocks. In fact, they only do so some of the time. The box corresponding to ‘John had just pushed him just before’ contains a blue block (*just before* is temporally deictic.) But in the ‘Max fell. John pushed him’ example, the ‘John pushed him’ box contains no blue blocks. If box  $\beta$  does not contain any blue blocks (there is no overt textual signal of time relative to context in the individual text span), then the tower cannot be built. The temporal relation in question is not entailed by the ‘box contents’, and thus the temporality of the relation passes the rhetorical criterion.

In my modification of the relational criterion, however, a LEGO tower of red and blue can be dismantled, and the component bricks can be placed into the respective boxes; the tower can then be reassembled (deduced) from the newly available components. The temporal relation in the ‘Max fell. John pushed him’ clause is pragmatically derived; it is not built up out of the box contents. But on closer inspection, when the LEGO tower corresponding to this relation drops out of the sky (is pragmatically inferred), it turns out to be nothing more than a tower of red and blue blocks. It can be taken apart, the blocks put into their respective sentence-boxes, and then reassembled (re-entailed) on the basis of what is now text-span-specific information. The temporal relation is, if not entailable, then ‘re-entailable’; so temporality fails this version of the relational criterion (as desired.) Causal relations, on the other hand, are purple; and no amount of dismantling LEGO towers will reduce purple towers to blue and red blocks. So causality passes the criterion.

## 7. Conclusion

RST is in crisis in the computational linguistics community. It was formulated as a rule-of-thumb, intuitive theory; but the lack of formal explicitness has made rhetorical analysts hold little confidence in its analyses—particularly as RST becomes widespread in text generation, where the formalisation of RST is an essential prerequisite. The responses to this crisis have ranged from Hovy’s *laissez-faire* stance (which does not address analysts’ intuitions that only certain types of relation belong in a rhetorical theory) to Korel’sky & Kittredge’s stratified RST (a completely unrecognisable variant of the original.)

Sanders *et al.*’s relational criterion is a formalisation of analysts’ intuitions as to what features are relevant to distinguishing between relations in a rhetorical theory. It is a sound basis for discussion, but it suffers in practice from a woolly statement of precisely how information is or is not local to a text span. Looked at carefully, it does not in fact reject temporality as a rhetorically relevant feature, although it claims to. In the light of Lascarides, Asher and Oberlander’s work on DICE (which clearly illustrates how real-world information can be inferred from a text using rhetorical and pragmatic information) I reformulate the relational criterion, such that rhetorically relevant parameters not only mustn’t be local to individual text spans, but should also be *irreducible* to information local to individual text spans.

The next logical step in this line of research (building on the feature-based analysis of RST in Nicholas (1994)) is to apply the relational criterion to the distinctions between relations in M&T’s inventory, and see how many of the relations pass the test of rhetorical distinctiveness. For example: M&T distinguish between VOLITIONAL CAUSE and NON-VOLITIONAL CAUSE; intuitively, this distinction seems semantic, and not rhetorically relevant. I would contend that VOLITIONAL CAUSE(*a, b*) can be reduced to CAUSE(*a, b*) and

*intend*( $Ag_a, b$ ): the event in *a* volitionally causes that in *b* iff *a* causes *b* and the agent of *a* intends *b*. NON-VOLITIONAL CAUSE, on the other hand, is reducible to CAUSE(*a*, *b*) and  $\neg$ *intend*( $Ag_a, b$ ). I would argue that *intend*( $Ag_a, b$ ) is information specific to *b*; therefore, the distinction between VOLITIONAL and NON-VOLITIONAL CAUSE is reducible to a distinction which follows deductively from segment-specific information, and is hence non-relational and not relevant to a rhetorical theory.<sup>11</sup>

Whether RST will survive its current crisis, particularly in computational linguistics, remains to be seen; in the interim, I hope this discussion helps to make the issues being debated somewhat more clear-cut.

## Appendix: Definitions of Rhetorical Relations

### CIRCUMSTANCE

<b>relation name:</b>	CIRCUMSTANCE
<b>constraints on N:</b>	none
<b>constraints on S:</b>	presents a situation (not unrealized)
<b>constraints on the N+S combination:</b>	S sets a framework in the subject matter within which R is intended to interpret the situation presented in N
<b>effect:</b>	R recognizes that the situation presented in S provides the framework for interpreting N
<b>locus of the effect:</b>	N and S

Example given: “[*N*] Probably the most extreme case of Visitors Fever I have ever witnessed was a few summers ago, [*S*] when I visited relatives in the Midwest.”

The satellite in a CIRCUMSTANCE relation sets a framework, e.g., a temporal or spatial framework, within which to interpret the nucleus. This function has been grammaticized in English in the form of circumstantial hypotactic clauses (M&T 1987:48)

### UNTIL

<b>relation name:</b>	UNTIL
<b>constraints on N:</b>	presents an action
<b>constraints on S:</b>	presents an unrealized situation
<b>constraints on the N+S combination:</b>	N is carried out as long as S is not yet true; S may result from N
<b>effect:</b>	R recognizes that N has to be carried out only as long as S is unrealized
<b>locus of the effect:</b>	N and S

Example given: “[*N*] Fill the radiator with new coolant [*S*] up to the filler opening.”

If we want to see these examples as made up from two minimal units at all [...] we should have a relation—named UNTIL or similarly—where the nucleus presents an action and the satellite presents a ‘stopping’ condition for the nucleus action, which often is at the same time the result of it.

In order to be helpful as a piece of advice, a ‘stopping condition’ in an UNTIL-relation must be easily observable. This is not the case in the following example:

“...leave the player for about an hour until the moisture evaporates”

One should not be misled by the proposition ‘until’ as surface cue; the example is better analyzed as VOLITIONAL RESULT relation (with the until-phrase as satellite), the ‘stopping condition’ is here given as temporal information (“for an hour”). (Rösner & Stede 1992:206)

<sup>11</sup>Volitionality is, of course, an important feature in human language. Then again, so is temporality; my argument is that neither lies in the proper domain of rhetorical theory.

## Symbols

The following symbols are used in this text:  $>$  is the defeasible conditional, ‘normally implies, implies by default.’  $\rightarrow$  is the material conditional.  $\leftrightarrow$  is material iff (if and only if.)  $\wedge$  is logical ‘and’, and  $\vee$  is logical ‘or’;  $\neg$  is logical ‘not’, and  $\exists$  is ‘there exists’.  $\leq$  means ‘temporally precedes.’  $\equiv$  means ‘equivalent to.’  $e_\alpha$  denotes the event described by text  $\alpha$ .  $A \models B$  means ‘B follows from A according to default logic.’  $[a, b]$  denotes an interval; a square bracket means that the endpoint is included in the interval, while a round bracket means it is excluded.  $\underline{x}$  is the order ‘Do  $x!$ ’, or the exhortation ‘ $x$  should be done!’ Semantic relations are italicised (e.g. *overlap*), to distinguish them from rhetorical relations, which appear in Small Caps (e.g. NARRATION).

## References

- 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.
- Grosz, B. J. & Sidner, C. L. 1986. Attention, Intentions and the Structure of Discourse. *Computational Linguistics* **12**:175–204.
- Halliday, M. & Hasan, R. 1976. *Cohesion in English*. London: Longman.
- Halliday, M. 1985. *An Introduction to Functional Grammar*. Baltimore: Edward Arnold.
- 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.
- Horacek, H. & Zock, M. (eds.) 1993. *New Concepts in Natural Language Preparation*. London: Pinter.
- Hovy, E. H. 1991. Parsimonious and Profligate Approaches to the Question of Discourse Structure Relations. In *Proceedings of the Workshop on Discourse Coherence*, University of Edinburgh, Edinburgh, 4–6 April. pp. 128–136.
- Knott, A. & Dale, R. 1993a. *Using Linguistic Phenomena to Motivate a Set of Rhetorical Relations*. Technical Report HCRC/RP–39, Human Communication Research Centre, University of Edinburgh.
- Knott, A. & Dale, R. 1993b. Choosing a Set of Rhetorical relations for Text Generation: A Data-Driven Approach. In *Proceedings of the Fourth European Workshop on Natural Language Generation*, Pisa, 28–30 April. pp. 163–166.
- Korelsky, T. & Kittredge, R. 1994. Towards Stratification of RST. In *Proceedings of the 32nd Annual Meeting of the Association for Computational Linguistics*, New Mexico State University, Las Cruces (NM), 27–30 June. pp. 52–55.
- Krifka-Dobeš, Z. & Novak, H.–J. 1993. From constituent planning to text planning. In Horacek & Zock. pp. 87–113.
- Lascardes, A. & Asher, N. 1991. Discourse Relations and Defeasible Knowledge. In *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, University of California, Berkeley (CA), 18–21 June. pp. 55–62.
- Lascardes, A., Asher, N. & Oberlander, J. 1992. Inferring Discourse Relations in Context. In *Proceedings of the 30th Annual Meeting of the Association for Computational Linguistics*, University of Delaware, Newark (DE), 28 June–2 July. pp. 1–8.
- Maier, E. & Hovy, E. 1993. Organising discourse structure relations using metafunctions. In Horacek & Zock. pp. 69–86.
- 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 (CA): Information Sciences Institute.
- 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.
- Moser, M. & Moore, J. 1993. Investigating Discourse Relations. In *Intentionality and Structure in Discourse Relations: Proceedings of a Workshop Sponsored by the Special Interest Group on Generation of the Association for Computational Linguistics*, Columbus (OH)
- Nicholas, N. 1994. *Problems in the Application of Rhetorical Structure Theory to Text Generation*. Unpublished Master’s thesis, University of Melbourne.



- Rösner, D. & Stede, M. 1992. Customising RST for the Automatic Production of Technical Manuals. In Dale *et al.* pp. 199–214.
- Sanders, T. & Spooren, W. 1990. Organizing the Set of Coherence Relations. In *Proceedings of the Fifth International Workshop on Natural Language Generation*, Dawson (PA)
- Sanders, T. J. M.; Spooren, W. P. M. & Noordman, L. G. M. 1992. Toward a Taxonomy of Coherence Relations. *Discourse Processes* **15**: 1–35.
- Sanders, T. 1993. On The Cognitive and Text-Analytic Status of Coherence Relations. In *Intentionality and Structure in Discourse Relations: Proceedings of a Workshop Sponsored by the Special Interest Group on Generation of the Association for Computational Linguistics*, Columbus (OH)
- Scott, D. R. & de Souza, C. S. 1990. Getting the Message Across in RST-based Text Generation. In Dale *et al.* pp. 47–73.
- Vander Linden, K. 1993. *Speaking of Actions: Choosing Rhetorical Status and Grammatical Form in Instructional Text Generation*. Ph.D. Dissertation. University of Colorado Tech Report CU–CS–654–93.