

Beyond elaboration: the interaction of relations and focus in coherent text

Alistair Knott Jon Oberlander Michael O'Donnell Chris Mellish

June 12, 2000

Abstract

This paper outlines a number of problems with RST's ELABORATION relation, and discusses a new model of text structure that results from leaving this relation out of the set of relations. In this model, trees of interclausal/intersentential relations account for the local coherence of a text, while its global coherence is accounted for by a separate device: global focus.

1 Introduction

Many theories of discourse propose that a coherent text is one whose clauses, sentences and text spans (or perhaps the propositions expressed by these text units) stand in particular relations to one another. The basic motivation in these theories stems from the observation that a text is more than a sequence of independent units: whether a particular unit makes sense in a given discourse depends not only on this unit by itself, but also on its relationship with the other units in the discourse. This claim has been spelled out in many different ways, but there are two requirements that any such theory must meet before it has empirical content and can be tested against the facts.

Firstly, a particular set of relations must be specified. It is vacuous to say that texts cohere in virtue of the relations that hold between their constituent units unless we specify what these relations are. There are as many 'possible relations between text units' as there are pairs of text units, and clearly since not all pairs of text units are coherent, we must select only some relations from this set. We can refer to the task of choosing a suitable set of relations as the task of developing a theory of **relation semantics**. Well-known theories of relation semantics include the set of 23 relations proposed by in the original formulation of RST (Mann and Thompson, 1988), the pair of relations DOMINANCE and SATISFACTION-PRECEDENCE proposed by Grosz and Sidner (1986) and the sets of conjunctive relations proposed by Halliday and Hasan (1976) and Martin (1983). Many of the papers in this volume are concerned with the task of defining a single class of relation, or of distinguishing a number of similar relations between one another.

Secondly, a theory of relations must provide an account of whereabouts in a coherent text relations are expected to be found. This account must begin by specifying what the atomic units of the analysis are. (Are they sentences? Clauses? Propositions within clauses? Units larger than sentences?) It must also state in a general way what structure of relations between these units will suffice to ensure its coherence. Clearly a text can be coherent without there being the right kind of relation between each pair of atomic units. Adjacency, or proximity, are

important factors. Often a notion of compositionality is also invoked, whereby two adjacent units linked by a relation are taken to form a new, composite unit, which can itself be linked by relations to other units. A theory which specifies whereabouts in a coherent text we can expect to find relations can be termed a theory of **span structure**. Many of the most influential theories of this kind (including RST and Grosz and Sidner’s theory) adopt the compositionality assumption in some form, and construe a coherent text as a tree of text units, in which complex units are formed from smaller units between which relations hold.

A theory of span structure and a theory of relation semantics are two logically separable components of a theory of discourse coherence. But naturally, adopting a specific theory of span structure can place constraints on what would be a sensible choice of theory of relation semantics. In this paper, we consider a case in point. The simple and parsimonious theory of span structure proposed by RST necessitates the inclusion in its theory of relation semantics of a rather idiosyncratic relation called OBJECT-ATTRIBUTE ELABORATION. We consider a number of problems with this relation in its own right, and also a number of problems with RST’s theory of span structure. We propose a revised account of relation semantics and span structure in which OBJECT-ATTRIBUTE ELABORATION is omitted, and which addresses these problems.

To illustrate both the problems with RST and the new account of discourse structure, we will use naturally-occurring and constructed texts in the genre of ‘museum guidebook descriptions’. The problems with RST were originally noticed when we built a text generation system that produces texts in this genre using a straightforward implementation of the theory. The theory implemented in the system was modified as a result, to overcome these problems.

2 RST’s theory of span structure

In RST’s theory of span structure, relations hold between **text spans**. Most relations have **nucleus-satellite structure**: one of the spans (the nucleus) is associated with the writer’s main communicative goal, and the other one (the satellite) is there to help bring about this goal, or to provide subsidiary information.

Atomic text spans are basically clauses. Complex text spans are structures called **schema applications**. A schema application for a nucleus-satellite relation is a set of adjacent text spans (either simple or complex), one of which is a nucleus, and the rest of which are linked to this nucleus by applications of a given nucleus-satellite relation. An example is given in Figure 1. *Nuc* is a nucleus span, *Sat1*, *Sat2* and *Sat3* are satellite spans linked to this nucleus by the relation *R*. The complex span *CS* is the complex span which is formed as a result.




Figure 1: A schema application of the nucleus-satellite relation *R*

In this paper, we will be considering three central assumptions underlying RST’s theory of span structure.

1. **Compositionality.** The first assumption relates to how the semantics of a complex text span are derived from the semantics of its constituent spans. The assumption is that a complex span comprising a nucleus and a number of satellites can be linked to another text span with a rhetorical relation iff its nucleus span can be so linked; in other words, for the purposes of linking spans together, the semantics of a span reduces to the semantics of its nucleus. This assumption is implicit in RST’s principal test for nuclearity, which specifies that the coherence of a text is largely preserved if the satellites in a given complex text span are removed, but lost if its nucleus is removed. The assumption has been stated more explicitly by Marcu (1997), who calls it the ‘strong compositionality’ assumption.
2. **Continuous constituency.** The second assumption relates to the distances over which relations are allowed to apply. Basically, RST requires that the nucleus N and satellite S of a relation R must either be adjacent text spans, or if not adjacent (as for instance in the case of Nuc and $Sat3$ in Figure 1), the text spans intervening between N and S must also be linked to N as satellites of the relation R .
3. **Tree structure.** In a coherent text, each text span (except for the complex span which constitutes the entire text) must be involved in exactly one schema application. This ensures firstly that there can be no sub-spans in the text that aren’t linked to any other spans, and secondly that there are no overlapping complex spans; basically, it specifies that a coherent text is a tree of schema applications.

These assumptions have proven very successful in identifying all and only well-structured texts. We illustrate with part of a text produced by ILEX-2, a generation system which delivers a sequence of descriptions of artefacts in a tour of a museum gallery.

- (1) (1) This jewel draws on natural themes for inspiration; (2) it is a remarkably fluid piece. (3) Indeed, Organic style jewels usually draw on natural themes for inspiration; (4) for instance the organic brooch we saw earlier looked crystalline.




Figure 2: RST Analysis of Example 1

The structure for this text is given in Figure 1. By the compositionality assumption, the top-level AMPLIFICATION relation holds between the complex spans (1–2) and (3–4) in virtue of their respective nuclear spans, (1) and (3). The expansions of (1) with (2), and of (3) with (4), take place independently of the higher-level relation. By continuous constituency, satellite spans appear adjacent to their nuclei. By tree structure, each sub-span in the text involved in exactly one schema application. In this case, adherence to these assumptions results in a well-structured text.

3 Some structural problems with ELABORATION

As well as being simple and parsimonious, RST’s theory of span structure is able to account for the coherence of a large number of texts. However, the theory has been criticised from several perspectives. For instance, the assumption of tree structure has been questioned by Sibun (1992), and the assumption of continuous constituency has been questioned by Kittredge *et al* (1991). Our central concern in this paper is to associate these structural problems with one RST relation in particular, namely OBJECT-ATTRIBUTE ELABORATION.¹

Mann and Thompson define this relation to hold between two spans if the nucleus ‘presents’ an object (i.e. contains a mention of it) and the satellite subsequently presents an attribute of that object. The precise meaning of ‘attribute’ is not clear, but any proposition which provides additional information about the object would seem to qualify. In the type of text which our system produces—a sequence of descriptions of a collection of related entities—this relation is heavily applicable, and the problems we note are thus quite widespread.

3.1 Discontinuous constituency

An initial problem is illustrated in the following text, taken from a museum guidebook.

- (2) (1) In the women’s quarters the business of running the household took place. (2) Much of the furniture was made up of chests arranged vertically in matching pairs (...). (3) Female guests were entertained in these rooms, which often had beautifully crafted wooden toilet boxes with fold-away mirrors and sewing boxes, and folding screens, painted with birds and flowers.
- (4) Chests were used for the storage of clothes...

In this text, an entity mentioned in the middle of the first paragraph, *chests*, becomes the central topic of the second paragraph. We can refer to this move pre-theoretically as a **resumption**.²

The move is clearly legitimate in the above context, and yet an analysis in terms of a tree of relations is difficult. The problem is that sentence (4) needs to be seen as the satellite of an ELABORATION relation, but the obvious nucleus for this relation—sentence (2)—is not accessible. If we analyse sentences (2) and (3) as ELABORATIONS of sentence 1, as seems necessary, we have effectively closed off sentence (2) as the nucleus for further ELABORATIONS. In order to treat sentence (4) as an ELABORATION of sentence (2), we would have to analyse sentence (3) as being subordinate to sentence (2): this analysis seems inappropriate; moreover, it makes the position of the paragraph break hard to explain. Note that we cannot just ignore the relationship between sentences (2) and (4) in our representation of the text: it is only because the chests are mentioned in the former sentence that they are a relevant topic for discussion. To account for coherence in this case, it seems we must either abandon compositionality, in some circumstances, or adopt a notion of discontinuous constituency for

¹Our objections do not extend to other types of ELABORATION; for instance, what Mann and Thompson call PROCESS-STEP ELABORATION or GENERALISATION-SPECIFIC ELABORATION. In what follows, references to ELABORATION are exclusively to the object-attribute variety, unless otherwise stated.

²The notion of a resumption bears some resemblance to Grosz and Sidner’s notion of a **digression**. This is a discourse segment which (a) is not related to the immediately preceding segment by dominance or satisfaction-precedence, and (b) contains mention of an entity salient in the interrupted segment. However, Grosz and Sidner’s definition implies that a link due to a common entity can only occur between *adjacent* segments; our claim is that resumptions can occur between non-adjacent segments.

text spans, or abandon the requirement that each subspan in a text is involved in at least one schema application.

A particularly common manifestation of this problem is in cases of parallelism within discourse structure. Especially in descriptive texts, it is common for a number of entities to be introduced sequentially in a sequence of spans, and then elaborated on in subsequent spans in the order of their introduction. Accounting for these subsequent mentions as ELABORATIONS of the spans where they were introduced is not possible without violating adjacency or compositionality constraints. Mann and Thompson acknowledge from the outset that RST cannot account for the constraints which apply in such contexts. McKeown (1985) deals extensively with cases of parallelism in text, although this account is not set in the context of a theory of coherence relations. Kittredge *et al* (1991) give several examples of parallelism; indeed, in one case they identify ELABORATION as the relation responsible for the problem.

3.2 Nuclearity and embedding

The preceding section presents a case where a ‘context-free’ theory of span structure under-generates the space of possible texts. There are also cases where it overgenerates; again, these relate principally to the ELABORATION relation.

There often seem to be difficulties in embedding ELABORATIONS within other relations. Consider this constructed text:

- (3) (1) Arts-and-Crafts jewels tend to be elaborate. (2) However, this jewel has a simple form.

This text contains a CONCESSION relation whose nucleus is (2) and whose satellite is (1). In principle, we could expand either span with additional relations. But note what happens when we embed an ELABORATION under span (1):

- (4) (1) Arts-and-Crafts jewels tend to be elaborate. (1a) They are often mass-produced.
(2) However, this jewel is simple in form.

Sentence (1a) elaborates on (1) by providing more information about Arts-and-Crafts jewels. However, it also makes it hard to attach sentence (2) to sentence (1). Note that there *is* a coherent interpretation of the text, if (1a) is treated as *expanding on the proposition* that Arts-and-Crafts jewels are elaborate, for instance by arguing for it, or by providing an example, rather than simply as ‘saying something else about Arts-and-Crafts jewels’. However, we have chosen the elaborating sentence to make these interpretations implausible. Besides, under these interpretations the embedded relation is no longer object-attribute ELABORATION; that is precisely our point.

Note also that the problem is not just due to difficulties with ‘high-level’ relations in general, or with ‘left-branching’ tree structures. Compare a text with different embedded relations:

- (5) (1) Arts-and-Crafts jewels tend to be elaborate. (1b) Ornateness was the fashion at the turn of the century. (1c) And not just in jewellery either. (2) However, this jewel is simple in form.

The structure of this text is given in Figure 3. In this text, there are two levels of embedding, not just one: sentence (1) is related to sentence (1b) via an EXPLANATION relation, and




Figure 3: Analysis of Example 5

sentence (1b) is itself related to sentence (1c) via an amplification relation. Nonetheless, the CONCESSION relation between sentences (1) and (2) is still intelligible. Note that there is no way that the relationship can be understood as applying between sentence (2) and the sentence immediately preceding it.

Arguably, there may be a limit to the depth of embedding permissible for any relation, particularly for left-branching RS trees. We will discuss this idea more in Section 6. But ELABORATION's apparent resistance to even the simplest kind of embedding suggests that it is qualitatively different from the other relations.

4 Problems with ELABORATION as a coherence relation

We now turn to some problems with the ELABORATION as a component of a theory of relation semantics.

4.1 ELABORATION as a relation between entities

One initial point to note is that the relation of ELABORATION is not really a *relation between propositions* in the same way that the other relations in RST are. A relation like EXPLANATION genuinely holds between two elements which are propositions: it holds if one proposition provides an explanation of the other, and there is no simpler way to state the relationship than this. It is not possible to identify subcomponents of the related propositions which stand in a relationship to each other that allows us to deduce that an EXPLANATION relation holds between the propositions. The same holds for the other RST relations. If a CAUSE relation holds between two propositions, is not possible to identify components of these propositions—for instance entities or predicates—whose relationship by itself allows us to deduce that a CAUSE relation holds between the propositions they are part of. But for the ELABORATION relation, this *is* possible. An ELABORATION relation between two propositions holds in virtue of a particular relationship (namely, identity) holding between component elements of the respective propositions (namely, entities). It is only *indirectly* a relationship between the propositions, in virtue of this direct relationship between entities. Many of the problems we will mention below seem to stem from this basic point.

4.2 Overlap with the focus metaphor

The discourse phenomena described by ELABORATION appear to overlap extensively with phenomena described by other theories of discourse, namely those concerned with focus structure.

Consider firstly theories of local focus; in particular, Grosz *et al's* (1995) account of centering. A primary concern for this theory is to catalogue the different discourse structures which can obtain in cases where two adjacent sentences make reference to a common entity. The issue is explored both in hypotheses about how this entity should be referred to in the second sentence (for example, pronominally) and about which sentence configurations make for 'good continuations'. The centering account is explicitly entity-based, and is expressed at a level of detail far greater than that given in the definition of ELABORATION, which *prima facie* covers the same cases. Moreover, it is not bound by the hierarchical constraints imposed on RST relations which were shown to be problematic for ELABORATION: adjacent sentences are related in chains, rather than in trees.

Consider also global focus. It is often useful to speak about the global focus of a passage of text, for instance if we are summarising it, or trying to resolve anaphora within it. But it is not possible to represent the global focus of a text within the vocabulary of RST. Consider a simple passage, in which an entity is described in a sequence of adjacent clauses. An RST analysis could identify the first of these clauses as the nucleus of an ELABORATION schema application, whose satellites are the remaining clauses. But this analysis accords a spurious significance to the *proposition* expressed by the first clause. It is not the proposition which is being elaborated on, but the entity.

Proponents of RST are likely to concede that notions of local and global focus are necessary in addition to the account that it provides. But our point is that when these extra primitives are included in a theory of text coherence, the ELABORATION relation essentially becomes redundant, and makes no contribution of its own. The aspects of text coherence which it represents are also modelled—and better modelled—by the entity-based metaphor of focus.

4.3 Linguistic signals

It has often been observed that ELABORATION is one of the few relations for which there are no conjunctive linguistic signals. There are simply no sentence or clause connectives for signalling this relation. Connectives like *indeed*, *in fact* or *also* do not always work: often, the best method of signalling this relation seems simply to be to close the nucleus sentence with a period, and begin a new sentence for the satellite.

Mann and Thompson are at pains not to tie relations directly to linguistic signals. But there would undeniably be advantages to being able to make such connections. In practice, computational treatments of RST, whether in text generation or discourse structure parsing, *do* link relations to surface signals. And ELABORATION is invariably treated differently from other relations in these contexts. For instance, in Scott and de Souza's (1990) list of distinctive methods for signalling RST relations in generated texts, the ELABORATION relation is to be signalled by a relative clause whose head noun denotes the entity being elaborated on. Marcu's (1997) algorithm for identifying the relations in a text from surface cues relies principally on discourse markers for all relations except ELABORATION (and JOINT); for these latter two relations, word co-occurrence measures provide the strongest surface indicators.

There are also theoretical reasons for holding that relations are associated with particular classes of linguistic expression. The present authors have argued that the set of linguistic resources available for signalling relations in a language can provide valuable evidence for determining how the set of relations in that language should be defined (Knott, 1996; Knott and Mellish, 1996), and that the lack of conjunctive signals for ELABORATION provides evidence that it is different from other relations.

5 An ELABORATIONless model of text coherence

While many of the problems with ELABORATION have been noted in the past, the question of what an account of discourse relations would look like without this relation has not been seriously considered. It is this question which we would like to address. In this section, we outline a revised version of RST in which ELABORATION is omitted from the set of relations. We would like to preserve as much as possible of the RST-based model, while taking account of the exceptions due to ELABORATION noted above.

We propose that the global coherence of a text is determined by global focus, rather than by a tree-structure of relations between high-level text spans.³ At a high level of structure, we take a coherent text to be a sequence of focus spaces which succeed each other in a legal manner.

We will term a focus space an **entity-chain**: basically a portion of text in which the global focus is some particular entity. An entity-chain is made up of a sequence of RS trees, each constructed just as in RST, but minus the ELABORATION relation. These trees can either be simple trees consisting of just one text span, or more complex trees with several layers of hierarchy. In each case, we can define the **top nucleus** of the tree to be the leaf-level text span which is reached by following the chain of nuclei from its root; in other words, it is the nucleus of the nucleus of (...) the nucleus of the tree. A legal entity-chain whose focus is entity E is one where the top nucleus of *each* tree is a fact about E .⁴ Note that the facts within a single tree do not all have to be about the entity in focus. Coherence between these facts is not determined by their having entities in common, but by there being relationships of the right sort between the propositions they express.

A legal sequence of entity-chains is a sequence in which the focussed entity in each chain is mentioned in a proposition within the n previous chains. In our text generation system, n is effectively set to 4, although it is likely that the value of n should vary depending on the length of intervening chains. Determining the factors which contribute to the value of n is a matter for further empirical investigation. Our main claim is that the admissibility of a chain with a particular focus at a particular point in a text is a function of its linear distance from the previous mention of the focussed entity, rather than of its relationship to the ‘right frontier’ of a discourse structure tree.

An example of a legal sequence of four entity-chains $EC1$, $EC2$, $EC3$ and $EC4$ is given in Figure 4. Within each entity-chain, atomic RS trees are denoted by rectangular boxes




Figure 4: A legal sequence of entity-chains

and non-atomic RS trees are denoted by triangles. The directed arcs indicate resumption relations: links from an entity-chain to the sentence which introduces it. Note that these arcs do not have to link adjacent entity-chains, and can cross one another.

The model of text structure just outlined has been implemented in the ILEX-2 text generation system; see Mellish *et al* (1998) for details. An example of a text generated by the

³In this respect, our proposal is similar to that made by Mooney *et al* (1990).

⁴A working definition of what it is for a fact to be ‘about’ a certain entity is given in Mellish *et al* (1998).

system is given below:

- (6) (1) This piece is a necklace. (2) It was designed by a jeweller called Jessie King. (3) It was designed in 1905. (4) It is made of silver and enamel.
- (5) Jessie King was a famous designer. (6) She was Scottish, (7) but she worked in London. (8) It was in London that this piece was made.
- (9) Like the previous piece, (10) this piece is in the Arts-and-Crafts style. (11) Although the previous piece had a simple shape, (12) Arts-and-Crafts style jewels tend to be elaborate; (13) for instance, this piece has detailed florals.

There are three entity-chains in the text: E_0 (spans 1–4) is about a particular jewel, E_1 (spans 5–8) is about the jewel’s designer, and E_2 (spans 9–13) is about the style it is in. Within these chains there are a number of local RS trees: spans 6–7 (top nucleus span 7), spans 9–10 (top nucleus span 10), and spans 11–13 (top nucleus span 12). Resumptions occur from (b) to (a), and from (c) to (a). Note that neither of these resumptions are to material in an adjacent text span. Nevertheless, the resulting text seems a good optimisation of focus and relation-based constraints.

6 Relations at higher levels of hierarchy

The idea of abolishing relations in the global structure of a text is certainly quite radical. It is a central tenet of RST that relations can apply between text spans of arbitrary size; and indeed, there are many complex texts in which relations do seem to apply at a high level of hierarchy. What should the present account say about these relations?

One thing to note immediately about these high-level relations is that they are not associated with surface conjunctive signals in the same way as low-level relations are. Low-level relations between clauses and sentences can typically be signalled directly by conjunctions, but conjunctions cannot be used to link arbitrarily large passages of text. If an explicit signal is needed, a slightly different mechanism is used, which involves a *new mention* of the top nucleus of the first span. Either this proposition is simply reiterated, in what Walker (1993, 1996) calls an ‘informationally-redundant utterance’, or the proposition is referred to *as an entity*, via the mechanism of nominalisation or discourse deixis (Webber, 1991). Assume we are given a large span of text S_1 , containing an argument that Kennedy was assassinated by the CIA, whose top nucleus is naturally enough the proposition *Kennedy was assassinated by the CIA*. If we want to continue with a second span S_2 concluding that we can’t trust the CIA, the three methods outlined above could be illustrated as follows.

- (7) (S1.) Given that *Kennedy was assassinated by the CIA*, the organisation is clearly untrustworthy. (Informationally-redundant utterance.)
- (8) (S1.) *Kennedy’s assassination by the CIA* proves that the organisation is untrustworthy. (Nominalisation.)
- (9) (S1.) *This* proves that the organisation is untrustworthy. (Discourse deixis—only possible when the top nucleus of S_1 can be referred to anaphorically.)

Using informationally-redundant utterances, any of the methods for signalling relations between small text spans are available for larger spans too. Likewise, the mechanisms of

nominalisation and discourse deixis provide the means for expressing high-level relations between propositions. But note that these latter methods for signalling high-level relations involve treating propositions as entities about which things can be predicated. What is more, nominalisation is a device which allows arbitrary reference to recent propositions in the text; the propositions which can be referred to are not limited to those on the right frontier of a discourse structure tree. Given these two considerations, we suggest that high-level relations signalled using nominalisations can be thought of, and are perhaps better thought of, in entity-based terms, as signals of *resumptions*. If we allow that any proposition in a text introduces *itself* as a possible topic for resumption, in addition to any entities it refers to directly, then the model of global text structure we presented in Section 5 seems to extend very well to the kind of high-level relations we have been discussing in this section.

7 Discussion

This paper has discussed a number of problems with RST's theories of span structure and relation semantics which stem from its use of the relation OBJECT-ATTRIBUTE ELABORATION. It argues that a better account of text coherence can be developed by abandoning this relation, and allowing that the metaphor of 'relations between propositions' only provides a *partial* account of text coherence. A new account of coherence is put forward in which a model of relations is supplemented with an entity-based model of focus structure. While previous accounts have suggested that relations and focus provide simultaneous constraints on coherence, the central idea in the new account is that the two adjacent text spans are coherent if *either* there is a suitable relation between the propositions they express, *or* they are linked by a legal focussing move.

There are three principal advantages of the new account. Firstly, by removing ELABORATION from the set of relations, we are able to eliminate some redundancy from any account of coherence which features constraints due both to relations and to local/global focus. Secondly, the new account promises to allow a tighter association between the primitives in a discourse theory and the linguistic means by which they are expressed: coherence relations can be associated with sentence and clause conjunctions, while focus-based moves are associated with nominal referring expressions. These associations are beneficial both for the first-order task of analysing texts, and the second-order task of defining the set of relation-based and entity-based primitives on which first-order analyses can draw. Thirdly, the division of labour between relations and focus in the new account produces a better match to the data in some respects. At low levels of hierarchy in a text, ELABORATION cannot be embedded inside RS trees in the way that other relations can. At high levels of hierarchy, non-local and crossing dependencies do seem to occur in text structure, but they seem to be restricted to cases of resumption, either of an entity mentioned in a recent proposition, or of this recent proposition itself.

The new proposal is certainly still at a preliminary stage of development. Empirical work is needed to investigate the claims about non-local and crossing resumption relations. The space of texts generated by our text planner provides some tentative evidence of the existence of text structures which RST cannot analyse, but a study of naturally-occurring text would provide a much better testbed for the theory. What is more, there remains much to be worked out in the new model. For one thing, the weaker constraints it imposes at the level of global structure may well lead it to overgenerate the space of coherent texts; additional constraints

may need to be specified. Additional constraints are also likely to be needed to determine the internal composition of entity-chains. These are avenues we are currently pursuing.

References

- Grosz, B. J. and Sidner, C. L. (1986). Attention, intentions, and the structure of discourse. *Computational Linguistics*, pages 175–203.
- Grosz, B. J., Joshi, A. K., and Weinstein, S. (1995). Centering: A framework for modeling the local coherence of discourse. *Computational Linguistics*, **21**(2), 203–225.
- Halliday, M. and Hasan, R. (1976). *Cohesion in English*. Longman.
- Kittredge, R., Korelsky, T., and Rambow, O. (1991). On the need for domain communication knowledge. *Computational Intelligence*, **7**, 305–314.
- Knott, A. (1996). *A Data-Driven Methodology for Motivating a Set of Coherence Relations*. Ph.D. thesis, Department of Artificial Intelligence, University of Edinburgh.
- Knott, A. and Mellish, C. (1996). A feature-based account of the relations signalled by sentence and clause connectives. *Language and Speech*, **39**(2–3), 143–183.
- Mann, W. C. and Thompson, S. A. (1988). Rhetorical structure theory: A theory of text organization. *Text*, **8**(3), 243–281.
- Marcu, D. (1997). *The rhetorical parsing, summarisation and generation of natural language texts*. Ph.D. thesis, Department of Computer Science, University of Toronto.
- Martin, J. R. (1983). Conjunction: The logic of English text. In J. S. Petöfi and E. Sözer, editors, *Micro and Macroconnectivity of Texts*, pages 1–72. Helmut Buske Verlag, Hamburg.
- McKeown, K. R. (1985). *Text Generation: Using Discourse Strategies and Focus Constraints to Generate Natural Language Text*. Cambridge University Press.
- Mellish, C., O'Donnell, M., Oberlander, J., and Knott, A. (1998). An architecture for opportunistic text generation. In *Proceedings of the ninth International Workshop on Natural Language Generation*, pages 28–37, Montréal.
- Mooney, D., Carberry, M. S., and McCoy, K. F. (1990). The generation of high-level structure for extended explanations. *COLING 90*, **2**, 276–281.
- Scott, D. R. and de Souza, C. S. (1990). Getting the message across in RST-based text generation. In R. Dale, C. Mellish, and M. Zock, editors, *Current Research in Natural Language Generation*. Academic Press.
- Sibun, P. (1992). Generating text without trees. *Computational Intelligence*, **8**(1), 102–122.
- Walker, M. (1993). *Informational redundancy and resource bounds in dialogue*. Ph.D. thesis, University of Pennsylvania.
- Walker, M. (1996). Limited attention and discourse structure. *Computational Linguistics*, **22**(2), 255–264.

Webber, B. (1991). Structure and ostension in the interpretation of discourse deixis. *Natural Language and Cognitive Processes*, **6**(2), 107–135.