

Discourse Structures for Text Generation

William C. Mann
USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292-6695

Abstract

Text generation programs need to be designed around a theory of text organization. This paper introduces *Rhetorical Structure Theory*, a theory of text structure in which each region of text has a central *nuclear* part and a number of *satellites* related to it. A natural text is analyzed as an example, the mechanisms of the theory are identified, and their formalization is discussed. In a comparison, Rhetorical Structure Theory is found to be more comprehensive and more informative about text function than the text organization parts of previous text generation systems.

1. The Text Organization Problem

Text generation is already established as a research area within computational linguistics. Although so far there have been only a few research computer programs that can generate text in a technically interesting way, text generation is recognized as having problems and accomplishments that are distinct from those of the rest of computational linguistics. Text generation involves creation of multisentential text without any direct use of people's linguistic skills; it is not computer-aided text creation.

Text planning is a major activity within text generation, one that strongly influences the effectiveness of generated text. Among the things that have been taken to be part of text planning, this paper focuses on just one: text organization. People commonly recognize that well-written text is organized, and that it succeeds partly by exhibiting its organization to the reader.

Computer generated text must be organized. To create

This research was supported by the Air Force Office of Scientific Research contract No. F49620-79-C-0181. The views and conclusions contained in this document are those of the author and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Office of Scientific Research of the U.S. Government.

text generators, we must first have a suitable theory of text organization. In order to be most useful in computational linguistics, we want a theory of text organization to have these attributes:

1. **comprehensiveness**: applicable to every kind of text;
2. **functionality**: informative in terms of how text achieves its effects for the writer;
3. **scale insensitivity**: applicable to every size of text, and capable of describing all of the various sized units of text organization that occur;
4. **definiteness**: susceptible to formalization and programming;
5. **generativity**: capable of use in text construction as well as text description.

Unfortunately, no such theory exists. Our approach to creating such a theory is described below, and then compared with previous work on text generation in Section 3.

2. Rhetorical Structure Theory

Creating a comprehensive theory of text organization is necessarily a very complex effort. In order to limit the immediate complexity of the task we have concentrated first on creating a descriptive theory, one which fits naturally occurring text. In the future the descriptive theory will be augmented in order to create a constructive theory, one which can be implemented for text generation. The term Rhetorical Structure Theory (RST) refers to the combination of the descriptive and constructive parts.

An organized text is one which is composed of discernible parts, with the parts arranged in a particular way and connected together to form a whole. Therefore a theory of text organization must tell at least:

1. What kinds of parts are there?
2. How can parts be arranged?
3. How can parts be connected together to form a whole text?

In RST we specify all of these jointly, identifying the organizational resources available to the writer.

2.1. Descriptive Rhetorical Structure Theory¹

What are the organizational resources available to the writer? Here we present the mechanisms and character of rhetorical structure theory by showing how we have applied it to a particular natural text. As each new construct is introduced in the example, its abstract content is described.

Our illustrative text is shown in Figure 2-1.²³ In the figure, we have divided the running text into numbered clause-like units.⁴

At the highest level, the text is a request addressed to CCC members to vote against making the nuclear freeze initiative (NFI) one of the issues about which CCC actively lobbies and promotes a position. The structure of the text at this level consists of two parts: the request (clause 13) and the material put forth to support the request (clauses 1 through 12).

2.1.1. The Request Schema --- 1-12; 13

To represent the highest level of structure, we use the *Request* schema shown in Figure 2-2. The *Request* schema is one of about 25 *schemas* in the current version of RST.

Each schema indicates how a particular unit of text structure is decomposed into other units. Such units are called *spans*. Spans are further differentiated into *text spans* and *conceptual spans*, text spans denoting the portion of explicit text being described, and conceptual spans denoting clusters of propositions concerning the subject matter (and sometimes the process of expressing it) being expressed by the text span.

¹The descriptive portion of rhetorical structure theory has been developed over the past two years by Sandra Thompson and me, with major contributions by Christian Matthiessen and Barbara Fox. They have also given helpful reactions to a previous draft of this paper.

²Quoted (with permission) from *The Insider*, California Common Cause state newsletter, 2.1, July 1982.

³We expect the generation of this sort of text to eventually become very important in Artificial Intelligence, because systems will have to establish the acceptability of their conclusions on heuristic grounds. AI systems will have to establish their credibility by arguing for it in English.

⁴Although we have not used technically-defined clauses as units, the character of the theory is not affected. The decision concerning what will be the finest-grain unit of description is rather arbitrary; here it is set by a preliminary syntax-oriented manual process which identifies low-level, relatively independent units to use in the discourse analysis. One reason for picking such units is that we intend to build a text generator in which most smaller units are organized by a programmed grammar [Mann & Matthiessen 3].

1. I don't believe that endorsing the Nuclear Freeze Initiative is the right step for California CC.
2. Tempting as it may be,
3. we shouldn't embrace every popular issue that comes along.
4. When we do so
5. we use precious, limited resources where other players with superior resources are already doing an adequate job.
6. Rather, I think we will be stronger and more effective
7. if we stick to those issues of governmental structure and process, broadly defined, that have formed the core of our agenda for years.
8. Open government, campaign finance reform, and fighting the influence of special interests and big money, these are our kinds of issues.
9. (New paragraph) Let's be clear:
10. I personally favor the initiative and ardently support disarmament negotiations to reduce the risk of war.
11. But I don't think endorsing a specific nuclear freeze proposal is appropriate for CCC.
12. We should limit our involvement in defense and weaponry to matters of process, such as exposing the weapons industry's influence on the political process.
13. Therefore, I urge you to vote against a CCC endorsement of the nuclear freeze initiative.

(signed) Michael Asimow, California Common Cause Vice-Chair and UCLA Law Professor

Figure 2-1: A text which urges an action

Each schema diagram has a vertical line indicating that one particular part is *nuclear*. The nuclear part is the one whose function most nearly represents the function of the text span analyzed in the structure by using the schema. In the example, clause 13 ("Therefore, I urge you to vote against a CCC endorsement of the nuclear freeze initiative.") is nuclear. It is a request. If it could plausibly have been successful by itself, something like clause 13 (without "Therefore") might have been used instead of the entire text. However, in this case, the writer did not expect that much to be enough, so some additional support was added.



Figure 2-2: The *Request* and *Evidence* schemas

The support, clauses 1 through 12, plays a *satellite* role in this application of the *Request* schema. Here, as in most cases, satellite text is used to make it more likely that the nuclear text will succeed. In this example, the writer is arguing that the requested action is right for the organization.

In Figure 2-2 the nucleus is connected to each satellite by a *relation*. In the text clause 13 is related to clauses 1 through 12 by a *motivation* relation. Clauses 1 through 12 are being used to motivate the reader to perform the action put forth in clause 13.

The relations relate the conceptual span of a nucleus with the conceptual span of a satellite. Since, in a text structure, each conceptual span corresponds to a text span, the relations may be more loosely spoken of as relating text spans as well.

The *Request* schema also contains an *enablement* relation. Text in an "enablement" relation to the nucleus conveys information (such as a password or telephone number) that makes the reader able to perform the requested action. In this example the option is not taken of having a satellite related to the nucleus by an "enablement" relation.

One or more schemas may be *instantiated* in a text. The pattern of instantiation of schemas in a text is called a *text structure*. So, for our example text, one part of its text structure says that the text span of the whole text corresponds to an instance of the *Request* schema, and that in that instance clause 13 is the text span corresponding to the schema nucleus and clauses 1 through 12 are the text span corresponding to a satellite related to the nucleus by a "motivation" relation.

In any instance of a schema in a text structure, the nucleus must be present, but all satellites are optional. We⁵ do not instantiate a schema unless it shows some decomposition of its text span, so at least one of the satellites must be present. Any of the relations of a schema may be instantiated indefinitely many times, producing indefinitely many satellites.

⁵Here and below, the knowledgeable person using RST to describe a text.

The schemas do not restrict the order of textual elements. There is a usual order, the one which is most frequent when the schema is used to describe a large text span; schemas are drawn with this order in the figures describing them apart from their instantiation in text structure. However, any order is allowed.

2.1.2. The *Evidence* Schema --- 1; 2-8; 9-12

At the second level of decomposition each of the two text spans of the first level must be accounted for. The final text span, clause 13, is a single unit. For more detailed description a suitable grammar (and other companion theories) could be employed at this point.

The initial span, clauses 1 through 12, consists of three parts: an assertion of a particular claim, clause 1, and two arguments supporting that claim, clauses 2 through 8 and 9 through 12. The claim says that it would not be right for CCC to endorse the nuclear freeze initiative (NFI). The first argument is about how to allocate CCC's resources, and the second argument is about the categories of issues that CCC is best able to address.

To represent this argument structure we use the *Evidence* schema, shown in Figure 2-2. Conceptual spans in an *evidence* relation stand as evidence that the conceptual span of the nucleus is correct.

Note that the *Evidence* schema could not have been instantiated in place of the *Request* schema as the most comprehensive structure of the text, because clause 13 urges an action rather than supporting credibility. The "motivation" relation and the "evidence" relation restrict the nucleus in different ways, and thus provide application conditions on the schemas. The relations are perhaps the most restrictive source of conditions on how the schemas may apply. In addition, there are other application conventions for the schema, described in Section 2.2.3.

The top two levels of structure of the text, the portion analyzed so far, are shown in Figure 2-3. The entire structure is shown in Figure 2-5.




Figure 2-3: The upper structure of the CCC text

At each level of structure it is possible to trace down the chain of nuclei to find a single clause which is representative of the entire level. Thus the representative of the whole text is clause 13 (about voting), the representative of the first argument is clause 6 (about being stronger and more effective), and the representative of the second argument is clause 12 (about limiting involvement to process issues).

2.1.3. The *Thesis/Antithesis* Schema --- 2-5; 6-8

The first argument is organized contrastively, in terms of one collection of ideas which the writer does not identify with, and a second collection of ideas which the writer does identify with. The first collection involves choosing issues on the basis of their popularity, a method which the writer opposes. The second collection concerns choosing issues of the kinds which have been successfully approached in the past, a method which the writer supports.

To account for this pattern we use the *Thesis/Antithesis* schema shown in Figure 2-4. The ideas the writer is rejecting, clauses 2 through 5, are connected to the nucleus (clauses 6 through 8) by a *Thesis/Antithesis* relation, which requires that the respective sections be in contrast and that the writer identify or not identify with them appropriately.

Notice that in our instantiations of the *Evidence* schema and the *Thesis/Antithesis* schema, the roles of the nuclei relative to the satellites are similar: Under favorable conditions, the satellites would not be needed, but under the conditions as the author conceives them, the satellites increase the likelihood that the nucleus will succeed. The assertion of clause 1 is more likely to succeed because the evidence is present; the antithesis idea is made clearer and more appealing by rejecting the competing thesis idea. The *Evidence* schema is different from the *Thesis/Antithesis* schema because evidence and theses provide different kinds of support for assertions.

2.1.4. The *Evidence* Schema --- 2-3; 4-5⁶

In RST, schemas are recursive. So, the *Evidence* schema can be instantiated to account for a text span identified by any schema, including the *Evidence* schema itself. This text illustrates this recursive character only twice, but mutual inclusion of schemas is actually used very frequently in general. It is the recursiveness of schemas which makes RST applicable at a wide range of scales, and which also allows it to describe structural units at a full range of sizes within a text.⁷

Clauses 2 and 3 make a statement about popular causes (centrally, that "we shouldn't embrace every popular issue that comes along"). Clauses 4 and 5 provide evidence that we shouldn't embrace them, in the form of an argument about effective use of resources.

The *Evidence* schema shown in Figure 2-2 has thus been used again, this time with only one satellite.

2.1.5. The *Concessive* Schema --- 2; 3

Clause 2 suggests that embracing every popular issue is tempting (and thus both attractive and defective). The attractiveness of the move is acknowledged in the notion of a popular issue. Clause 3 identifies the defect: resources are used badly.

The corresponding schema is the *Concessive* schema, shown in Figure 2-4. The *concession* relation relates the conceded conceptual span to the conceptual span which the writer is emphasizing. The "concession" relation differs from the "thesis/antithesis" relation in acknowledging the conceptual

⁶Except for single-clause text spans, the structure of the text is presented depth-first, left to right, and shown in Figure 2-5.

⁷This contrasts with some approaches to text structure which do not provide structure between the whole-text level and the clause level. Stories, problem-solution texts, advertisements, and interactive discourse have been analyzed in that way.

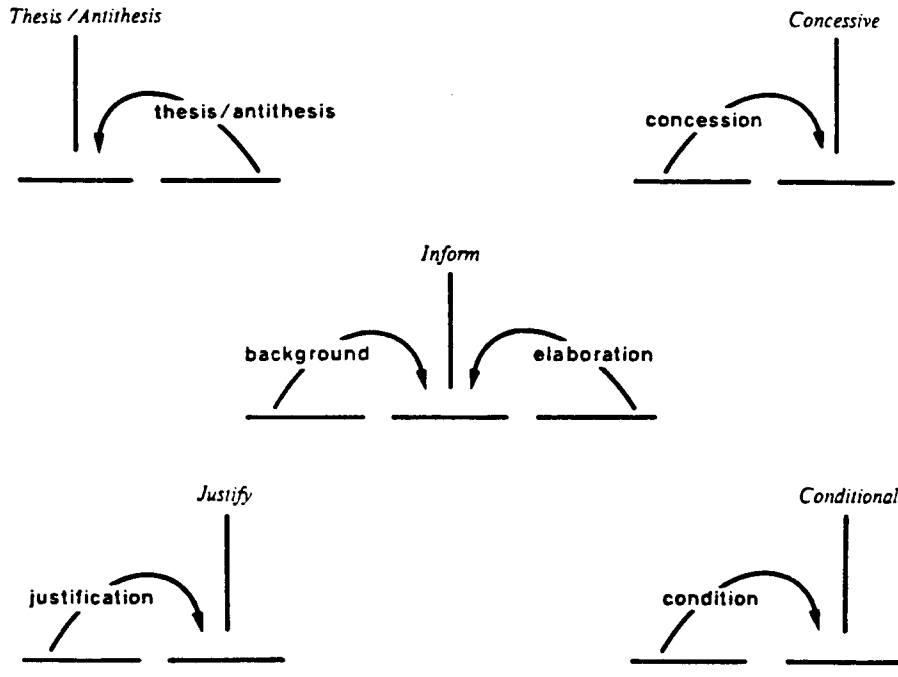


Figure 2-4: Five other schemas




Figure 2-5: The full rhetorical structure of the CCC text

span of the satellite. The strategy for using a concessive is to acknowledge some potential detractor or refutation of the point to be made. By accepting it, it is seen as not contradictory with other beliefs held in the same context, and thus not a real refutation for the main point.

Concessive structures are abundant in text that argues points which the writer sees as unpopular or in conflict with the audience's strongly held beliefs. In this text (which has two *Concessive* structures), we can infer that the writer believes that his audience strongly supports the NFI.

2.1.6. The *Conditional* Schema --- 4; 5

Clauses 4 and 5 present a consequence of embracing "every popular issue that comes along." Clause 4 ("when we do so") presents a condition, and clause 5 a result (use of resources) that occurs specifically under that condition. To express this, we use the *Conditional* schema shown in Figure 2-4. The condition is related to the nuclear part by a *condition* relation, which carries the appropriate application restrictions to maintain the conditionality of the schema.

2.1.7. The *Inform* Schema --- 6-7; 8

The central assertion of the first argument, in clauses 6 through 8, is that CCC can be stronger and more effective under the condition that it sticks to certain kinds of issues (implicitly excluding NFI). This assertion is then elaborated by exemplifying the kinds of issues meant.

This presentation is described by applying the *Inform* schema shown in Figure 2-4. The central assertion is nuclear, and the detailed identification of kinds of issues is related to it by an *elaboration* relation. The option of having a span in the instantiation of the *Inform* schema related to the nucleus by a *background* relation is not taken.

This text is anomalous among expository texts in not making much use of the *Inform* schema.⁸ It is widely used, in part because it carries the "elaboration" relation. The "elaboration" relation is particularly versatile. It supplements the nuclear statement with various kinds of detail, including relationships of:

1. set:member
2. abstraction:instance
3. whole:part
4. process:step
5. object:attribute

⁸It is also anomalous in another way: the widely used pattern of presenting a problem and its solution does not occur in this text.

2.1.8. The *Conditional* Schema --- 6; 7

This second use of the *Conditional* schema is unusual principally because the condition (clause 7) is expressed after the consequence (clause 6). This may make the consequence more prominent or make it seem less uncertain.

2.1.9. The *Justify* Schema --- 9; 10-12

The writer has argued his case to a conclusion, and now wants to argue for this unpopular conclusion again. To gain acceptance for this tactic, and perhaps to show that a second argument is beginning, he says "Let's be clear." This is an instance of the *Justify* schema, shown in Figure 2-4. Here the satellite is attempting to make acceptable the act of expressing the nuclear conceptual span.

2.1.10. The *Concessive* Schema --- 10; 11-12

The writer again employs the concessive schema, this time to show that favoring the NFI is consistent with voting against having CCC endorse it. In clause 10, the writer concedes that he personally favors the NFI.

2.1.11. The *Thesis/Antithesis* Schema --- 11; 12

The writer states his position by contrasting two actions: CCC endorsing the NFI, which he does not approve, and CCC acting on matters of process, which he does approve.

2.2. The Mechanisms of Descriptive RST

In the preceding example we have seen how rhetorical schemas can be used to describe text. This section describes the three basic mechanisms of descriptive RST which have been exemplified above:

1. Schemas
2. Relation Definitions
3. Schema Application Conventions

2.2.1. Schemas

A schema is defined entirely by identifying the set of relations which can relate a satellite to the nucleus.

2.2.2. Relation Definitions

A relation is defined by specifying three kinds of information:

1. A characterization of the nucleus,
2. A characterization of the satellite,
3. A characterization of what sorts of interactions between the conceptual span of the nucleus and the conceptual span of the satellite must be plausible.⁹

⁹All of these characterizations must be made properly relative to the writer's viewpoint and knowledge.

In addition, the relations are heavily involved in implicit communication; if this aspect is to be described, the relation definition must be extended accordingly. This aspect is outside of the scope of this paper but is discussed at length in [Mann & Thompson 83].

So, for example, to define the "motivation" relation, we would include at least the following material:

1. The nucleus is an action performable but not yet performed by the reader.
2. The satellite describes the action, the situation in which the action takes place, or the result of the action, in ways which help the reader to associate value assessments with the action.
3. The value assessments are positive (to lead the reader to want to perform the action).

2.2.3. Schema Application Conventions

Most of the schema application conventions have already been mentioned:

1. One schema is instantiated to describe the entire text.
2. Schemas are instantiated to describe the text spans produced in instantiating other schemas.
3. The schemas do not constrain the order of nucleus or satellites in the text span in which the schema is instantiated.
4. All satellites are optional.
5. At least one satellite must occur.
6. A relation which is part of a schema may be instantiated indefinitely many times in the instantiation of that schema.
7. The nucleus and satellites do not necessarily correspond to a single uninterrupted text span.

Of course, there are strong patterns in the use of schemas in text: relations tend to be used just once, nucleus and satellites tend to occur in certain orders, and schemas tend to be used on uninterrupted spans of text.

The theory currently contains about 25 schemas and 30 relations.¹⁰ We have applied it to a diverse collection of approximately 100 short natural texts, including administrative memos, advertisements, personal letters, newspaper articles, and magazine articles. These analyses have identified the usual patterns of schema use, along with many interesting exceptions.

The theory is currently informal. Applying it requires making judgments about the applicability of the relations, e.g., what counts as evidence or as an attempt to motivate or justify some action. These are complex judgments, not easily formalized.

¹⁰In this paper we do not separate the theory into framework and schemas, although for other purposes there is a clear advantage and possibility of doing so.

In its informal form the theory is still quite useful as a part of a linguistic approach to discourse. We do not expect to formalize it before going on to create a constructive theory. (Of course, since the constructive theory specifies text construction rather than describing natural texts, it need not depend on human judgements in the same way that the descriptive theory does.)

2.3. Assessing Descriptive RST

The most basic requirement on descriptive RST is that it be capable of describing the discernible organizational properties of natural texts, i.e., that it be a theory of discourse organization. The example above and our analyses of other texts have satisfied us that this is the case.¹¹

In addition, we want the theory to have the attributes mentioned in Section 1. Of these, descriptive RST already satisfies the first three to a significant degree:

1. **comprehensiveness:** It has fit many different kinds of text, and has not failed to fit any kind of non-literary monologue we have tried to analyze.
2. **functionality:** By means of the relation definitions, the theory says a great deal about what the text is doing for the writer (motivating, providing evidence, etc.).
3. **scale insensitivity:** The recursiveness of schemas allows us to posit structural units at many scales between the clause and the whole text. Analysis of complete magazine articles indicates that the theory scales up well from the smaller texts on which it was originally developed.

We see no immediate possibility of formalizing and programming the descriptive theory to create a programmed text analyzer. To do so would require reconciling it with mutually compatible formal theories of speech acts, lexical semantics, grammar, human inference, and social relationships, a collection which does not yet exist. Fortunately, however, this does not impede the development of a constructive version of RST for text generation.

2.4. Developing a Constructive RST

Why do we expect to be able to augment RST so that it is a formalizable and programmable theoretical framework for generating text? Text appears as it does because of intentional activity by the writer. It exists to serve the writer's purposes. Many

¹¹In another paper, we have shown that implicit communication arises from the use of the relations, that this communication is specific to each relation, and that as linguistic phenomena the relations and their implicit communication are not accounted for by particular existing discourse theories [Mann & Thompson 83].

of the linguistic resources of natural languages are associated with particular kinds of purposes which they serve: questions for obtaining information, marked syntactic constructions for creating emphasis, and so forth. At the schema level as well, it is easy to associate particular schemas with the effects that they tend to produce: the *Request* schema for inducing actions, the *Evidence* schema for making claims credible, the *Inform* schema for causing the reader to know particular information, and so forth. Our knowledge of language in general and rhetorical structures in particular can be organized around the kinds of human goals that the linguistic resources tend to advance.

The mechanisms of RST can thus be described within a more general theory of action, one which recognizes means and ends. Text generation can be treated as a variety of goal pursuit. Schemas are a kind of means, their effects are a kind of ends, and the restrictions created by the use of particular relations are a kind of precondition to using a particular means.

Goal pursuit methods are well precedented in artificial intelligence, in both linguistic and nonlinguistic domains [Appelt 81, Allen 78, Cohen 78, Cohen & Perrault 77, Perrault & Cohen 78, Cohen & Perrault 79, Newell & Simon 72]. We expect to be able to create the constructive part of RST by mapping the existing part of RST onto AI goal pursuit methods. In particular computational domains, it is often easy to locate formal correlates for the notions of evidence, elaboration, condition, and so forth, that are expressed in rhetorical structure; the problem of formalization is not necessarily hard.

At another level, we have some experience in using RST informally as a writer's guide. This paper and others have been written by first designing their rhetorical structure in response to stated goals. For this kind of construction, the theory seems to facilitate rather than impede creating the text.

3. Comparing RST to Other Text Generation Research

Given the mechanisms and example above, we can compare RST to other computational linguistic work on text generation.¹² The most relevant and well known efforts are by Appelt (the KAMP system [Appelt 81]), Davey (the PROTEUS system [Davey 79]), Mann and Moore (the KDS system [Mann & Moore 80, Mann & Moore 81]), McDonald (the MUMBLE system

[McDonald 80]) and McKeown (the TEXT system [McKeown 82]). All of these are informative in other areas but, except for McKeown, they say very little about text organization.

Appelt acknowledges the need for a discourse component, but his system operates only at the level of single utterances. Davey's excellent system uses a simple fixed narrative text organization for describing tic-tac-toe games: moves are described in the sequence in which they occurred, and opportunities not taken are described just before the actual move which occurred instead. Mann and Moore's KDS system organizes the text, but only at the whole-text and single-utterance levels. It has no recursion in text structure, and no notion of text structure components which themselves have text structure. McDonald took as his target what he called "immediate mode," attempting to simulate spontaneous unplanned speech. His system thus represents a speaker who continually works to identify something useful to say next, and having said it, recycles. It operates without following any particular theory of text structure and without trying to solve a text organization problem.

McKeown's TEXT system is the only one of this collection that has any hint of a scale-insensitive view of text structure. It has four programmed "schemas" (limited to four mainly by the computational environment and task). Schemas are defined in terms of a sequence of text regions, each of which satisfies a particular "rhetorical predicate." The sequence notation specifies optionality, repeatability, and allowable alternations separately for each sequence element. Recursion is provided by associating schemas with particular predicates and allowing segments of text satisfying those predicates to be expressed using entire schemas. Since there are many more predicates than schemas, the system as a whole is only partially recursive.

McKeown's approach differs from RST in several ways:

1. McKeown's schemas are ordered, those of RST unordered.
2. Repetition and optionality are specified locally; in RST they are specified by a general convention.
3. McKeown's schemas do not have a notion of a nuclear element.
4. McKeown has no direct correlate of the RST relation. Some schema elements are implicitly relational (e.g., an "attributive" element must express an attribute of something, but that thing is not located as a schema element). The difference is reduced by McKeown's direct incorporation of "focus."

The presence of nuclear elements in RST and its diverse collection of schemas make it more informative about the functioning of the texts it describes. Its relations make the

¹²Relating RST to the relevant *linguistic* literature is partly done in [Mann & Thompson 83], and is outside the scope of this paper. However, we have been particularly influenced by Grimes [Grimes 75], Hobbs [Hobbs 76], and the work of McKeown discussed below.

connectivity of the text more explicit and contribute strongly to an account of implicit communication.

Beyond these differences, McKeown's schemas give the impression of defining a more finely divided set of distinctions over a narrower range. The four schemas of TEXT seem to cover a range included within that of the RST *Inform* schema, which relies strongly on its five variants of the "elaboration" relation. Thus RST is more comprehensive, but possibly coarser-grained in providing varieties of description.

Our role for text organization is also different from McKeown's. In the TEXT system, the text was organized by a schema-controlled search over things that are permissible to say. In constructive RST, text will be organized by goal pursuit, i.e., by goal-based selection. For McKeown's task the difference might not have been important, but the theoretical differences are large. They project very different roles for the writer, and very different top-level general statements about the nature of text.

Relative to all of these prior efforts, RST offers a more comprehensive basis for text organization. Its treatment of order, optionality, organization around a nucleus, and the relations between parts are all distinct from previous text generation work, and all appear to have advantages.

4. Summary

A text generation process must be designed around a theory of text organization. Most of the prior computational linguistic work offers very little content for such a theory. In this paper we have described a new theoretical approach to text organization, one which is more comprehensive than previous approaches. It identifies particular structures with particular ways in which the text writer is served. The existing descriptive version of the theory appears to be directly extendible for use in text construction.

References

- [Allen 78] Allen, J., *Recognizing Intention in Dialogue*, Ph.D. thesis, University of Toronto, 1978.
- [Appelt 81] Appelt, D., Planning natural language utterances to satisfy multiple goals. Forthcoming Ph.D. thesis, Stanford University.
- [Cohen 78] Cohen, P. R., *On Knowing What to Say: Planning Speech Acts*, University of Toronto, Technical Report 118, 1978.
- [Cohen & Perrault 77] Cohen, P. R., and C. R. Perrault, "Overview of 'planning speech acts'," in *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, Massachusetts Institute of Technology, August 1977.
- [Cohen & Perrault 79] Cohen, P. R., and C. R. Perrault, "Elements of a plan-based theory of speech acts," *Cognitive Science* 3, 1979.
- [Davey 79] Davey, A., *Discourse Production*, Edinburgh University Press, Edinburgh, 1979.
- [Grimes 75] Grimes, J. E., *The Thread of Discourse*, Mouton, The Hague, 1975.
- [Hobbs 76] Hobbs, J., *A Computational Approach to Discourse Analysis*, Department of Computer Science, City College, City University of New York, Technical Report 76-2, December 1976.
- [Mann & Matthiessen 3.] Mann, W. C., and C. M. I. M. Matthiessen, *Nigel: A Systemic Grammar for Text Generation*, USC/Information Sciences Institute, RR-83-105, February 1983. The papers in this report will also appear in a forthcoming volume of the *Advances in Discourse Processes Series*, R. Freedle (ed.): *Systemic Perspectives on Discourse: Selected Theoretical Papers from the 9th International Systemic Workshop*, to be published by Ablex.
- [Mann & Moore 80] Mann, W. C., and J. A. Moore, *Computer as Author--Results and Prospects*, USC/Information Sciences Institute, RR-79-82, 1980.
- [Mann & Moore 81] Mann, W. C., and J. A. Moore, "Computer generation of multiparagraph English text," *American Journal of Computational Linguistics* 7, (1), January - March 1981.
- [Mann & Thompson 83] Mann, W. C., and S. A. Thompson, *Relational Propositions in Discourse*, USC/Information Sciences Institute, Marina del Rey, CA 90291, Technical Report RR-83-115, July 1983.
- [McDonald 80] McDonald, David D., *Natural Language Production as a Process of Decision-making under Constraints*, Ph.D. thesis, MIT, Cambridge, Mass., November 1980.
- [McKeown 82] McKeown, K.R., *Generating Natural Language Text in Response to Questions about Database Structure*, Ph.D. thesis, University of Pennsylvania, 1982.
- [Newell & Simon 72] Newell, A., and H. A. Simon, *Human Problem Solving*, Prentice-Hall, Englewood Cliffs, N.J., 1972.
- [Perrault & Cohen 78] Perrault, C. R., and P. R. Cohen, *Planning Speech Acts*, University of Toronto, Department of Computer Science, Technical Report, 1978.