

The Local Organization of Text

Penelope Sibun

Department of Computer and Information Science

University of Massachusetts

Amherst, MA 01003

sibun@cs.umass.edu or penni@umass.bitnet

Abstract

In this paper, I present a model of the local organization of extended text. I show that texts with weak rhetorical structure and strong domain structure, such as descriptions of houses, digital circuits, and families, are best analyzed in terms of local domain structure, and argue that global structures that may be inferred from a domain are not always appropriate for constructing descriptions in the domain. I present a system I am implementing that uses *short-range strategies* to organize text, and show how part of a description is organized by these strategies. I also briefly discuss a model of incremental text generation that dovetails with the model of local organization presented here.

Motivation for local organization

The approach to organizing extended text described here has both psychological and computational motivation. It aims both to model how people use language and to provide a flexible architecture for a system's language use. In this section, I describe the empirical data that form the basis of this research, and characterize the local organization of the collected texts. In the next two sections, I describe a computational architecture to implement local text organization and discuss its advantages of generality and flexibility, and give an example of how this architecture works.

An extended text has a structure; this structure is a description of how the components relate so that sense can be made of the whole. Two sources of this organization are *rhetorical structure*, which describes the way elements of the text fit together, and *domain structure*, which describes relations among domain objects. For this research I chose three domains with strong domain structure, and a task—description—with weak rhetorical structure. I have tape-recorded 29 people giving descriptions of house layouts, electronic circuit layouts, and family relationships. Description fragments of a house and of a family, and the questions asked to obtain the descriptions, are given in figure 1. (Because of space considerations, the fragments are somewhat abbreviated.)

Many approaches to text organization¹ are based on analyses of text in terms of rhetorical structure. However, there are few segments of text with interesting rhetorical structure in my corpus. For example, an analysis of the texts using Mann and Thompson's (1987) Rhetorical Structure Theory (RST) would result primarily in the relations sequence and joint and would contain few of the the relations like evidence or justify that give RST its descriptive power. Similarly, it is unclear what work a system like that of Grosz and Sidner (1986) would do in analyzing a description.

Since the structure of descriptions cannot be analyzed adequately with rhetorical relations, perhaps it can be explained in terms of the domain. Houses, chips, and families are strongly structured. A family's relationships can be captured in a family tree; one might suppose that a description of the family would also be organized in this way. A house can be encoded in a number of ways; for instance, it has a component hierarchy, being composed of rooms composed of furnishings. Linde (1974) has proposed another comprehensive structure for houses: a phrase structure grammar that determines how the rooms may be visited in a traversal of a house layout.

Surprisingly, these global, hierarchical domain structures are not exploited in the organization of descriptions in my corpus. While family trees and composition hierarchies can be inferred from descriptions of families and houses, this does not mean that these structures guide the process of organizing them. For instance, my family informants did not simply construct their descriptions by starting at the root of the appropriate tree and doing a depth-first or breadth-first traversal of it. Instead, to select a next family member to talk about, they would apply one of several criteria. Generally, a sibling, spouse, parent, or child would be the next choice, and this might incidentally constitute part of a tree walk. But that this choice is local is evidenced by the next choice, which may not be construable, in any

¹What I call text organization is usually referred to as text planning.

on our righthand side would be a door,
 which leads to Penni's room
 and you walk in there...
 and there are two windows,
 in the...opposite corner from the one in
 which you enter
 one's on the lefthand wall,
 and one's on the wall that you would be
 facing
 then, on the righthand side...of her room,
 is the closet

Fragment of a house description.

In response to the question: "Could you please describe for me the layout of your house."

there's my mother Katharine, my father John,
 my sister Penni, and me
 it's my mother's relatives that we go to see
 Margaret and Bill,
 who are Mommy's...
 um...Margaret's my great-aunt
 so it must be Mommy's aunt

Fragment of a family description.

In response to the question: "Can you tell me how everyone who comes to Thanksgiving is related to each other?"

Figure 1: Sample descriptions.

principled way, as part of the overall structure of the description that one might have postulated at the previous step. Where to begin the family description also appears to be a locally conditioned choice. Informants begin at various points, such as themselves or long-dead progenitrixes, but the majority start their description of the family by mentioning the hostess and host of the Thanksgiving dinner they are attending; we may suppose that at a different time of year the descriptions are likely to start off differently.

Further evidence that people do not structure their descriptions using obvious global domain structures may be adduced from examples in which speakers explicitly deny knowledge of such structures, as in the following fragment.

and also...um...
 Eleanor and Elizabeth come
 who are...cousins of...all of us...um...
 I don't know what generation cousins they
 are

Here, the speaker shows by her description of two family members that she does not know her relationship to them, even though it would be clear if a family tree were being used to organize the description (the women in question are in fact first cousins twice removed of the speaker).

Genealogical trees, phrase structure grammars, and component hierarchies are useful for succinctly representing information about houses, chips, and families

But there is no *a priori* reason to suppose that a description of such things are the products of such easily articulable schemas or grammars. When we examine texts of the sort that we wish to generate, we must distinguish the mechanisms that direct the process of choosing what to say next from a retrospective description of its result.

The texts I have collected can be best analyzed as *locally organized* by a process of deciding *what to say next*. This decision is based principally on what has already been said and what is currently available to say. For example, if one has just mentioned a large window in the kitchen, one can mention whatever is to the left of the window, whatever is to the right of it, which way it is facing, what it looks like, or how it is similar to a window in another room of the house. If one has mentioned an aunt, one can give her name, say whether she is married, mention her sister, enumerate her children, or talk about how much money she earns.

The strong domain structure of subjects like houses, chips, and families ensures that a description can be continued from any point: once a description has been started, there is always something, often many things, that can be said next. In structured domains, there is always a default choice for the next thing to say. In spatial domains like houses, spatial proximity provides this default. Everything in a house, be it a room, a wall, or a kitchen appliance, is next to something else. Spatial proximity does not *constrain* house descriptions,

but it ensures that a description does not come to a premature dead end.

Descriptions are finite. Though there may always be more to say, there are points at which a description may stop, when the task may be considered accomplished. Linde (1974) proposed a *completeness criterion* for house descriptions, which is reflected in my data as well as hers. It states that a description may stop any time after all the rooms have been mentioned, but it is not complete if it stops before. A similar criterion holds in the family descriptions collected: they were given in answer to the question, "Can you tell me how everyone who comes to Thanksgiving is related to each other?" In this case, then, the criterion is mentioning everyone who attends.

Knowing how to continue and knowing when to stop together ensure that a description can be generated depending solely on local organization. The strong domain structure of houses and families makes the working of these mechanisms for continuation and termination particularly clear, and thus these domains are a good site for studying this approach to organizing text. However, the local organization of text is also evident in many other uses of language. People's conversation is often locally organized (Levinson, 1983); some interactive systems are currently being designed with this approach (Frohlich & Luff, 1989).

Because I am interested not only in how a program may organize text but also in how people do so, I study people speaking rather than people writing. A written text may be edited and reorganized, and this process often involves explicitly thinking about rhetorical structure. A spoken description is likely to require that the speaker organize her text locally—she cannot plan it out ahead of time. Studying spoken text reveals more of the underlying mechanisms of language, because time constraints and the inability to edit what has already been said make post-processing impossible.

Computational architecture

I am implementing a system that employs local organization of text as described in this paper. The implementation comprises: a semantic net knowledge base; an organizer composed of strategies and metastrategies; and a generator. Local organization is achieved using *short-range strategies*, each of which is responsible for organizing only a short segment of text, between a word and a clause in length.

Until recently, I have employed Mumble-86 (Meter et al., 1987) as the generator for this system. Construction is underway, however, on a simpler generator that more accurately implements the principles of generation implied by the structure of the organizer. The system is currently implemented only for descriptions of houses; the examples in this section will thus be drawn from the house domain.

The knowledge base is a semantic net that encodes

the objects, properties, and relations of a domain. The organizer keeps a pointer to the current node in the description. A strategy describes the current node and others related to it via local connections in the network. Strategies are selected sequentially, based on local conditions; if there is no strategy available or if more than one strategy is appropriate, *metastrategies* (Davis, 1980) resolve the conflict, using a technique similar to universal subgoaling in Soar (Laird, Newell & Rosenbloom, 1987). Metastrategies are responsible for control: they sequence and combine strategies. Like strategies, they can cause the production of text.

This architecture has several advantages. First, it is flexible: because there is no fixed priority scheme and because strategies are small, the strategies can be combined in a variety of ways to produce texts that are constrained only by the appropriateness of each strategy as determined by the strategies' interaction with the knowledge base. Second, the architecture is extensible: new strategies can easily be added to extend the organizer to different types of text. Finally, the organizer is mainly domain-independent: while some strategies may be particular to houses, most strategies are not.

The strategies are applied to the knowledge base and select the items that make up the description. The strategies find the appropriate lexical items(s) for each knowledge base item that is expressed; these lexical items and the knowledge base items themselves are determined by the domain. While the strategies are simple, complex behavior emerges from their interaction with the knowledge base; this locally organizes the extended text.

Each strategy falls into one of four classes, with varying degrees of domain independence: *discourse cue strategies*; *linguistic strategies*; *parameterizable domain-independent strategies*; and *semi-domain-independent strategies*. Figure 2 gives examples of each.

Of the domain-independent strategies, discourse cue strategies focus attention, in a way similar to the clue words described by Reichman (1985), and linguistic strategies mention objects and associated properties. *mention-salient-object*, used to say "there is a window," may as easily express "there is a penguin" or "there is a policy." *describe-object* is similarly all-purpose, and can produce "the window with two flanking windows" or "the man with one black shoe."

The parameterizable domain-independent strategies have slightly different textual realizations in different domains, but these differences can be captured by parameters. The example given in figure 2 is typical: the strategy is realized as a prepositional phrase in each domain, and only the preposition changes.

The semi-domain-independent strategies accomplish tasks such as a sweep that seem particular to the domain, but are similar to tasks in other domains. A sweep begins at an object and names another bearing some spatial relationship to it, and then another object

Discourse Cue Strategies

introduce/shift-attention	“then,” “and,” “now”
refer-back/reinforce	“again,” “once more”

Linguistic Strategies

mention-salient-object	“there is <i>x</i> ” “we have <i>x</i> ”
describe-object	“the <i>x</i> is <i>y</i> ,” “the <i>y</i> <i>x</i> ” “the <i>x</i> (which) has <i>z</i> ” “the <i>x</i> with <i>z</i> ”

Parameterizable Domain-independent Strategies

situate	“in the kitchen” “during the morning” “about the election”
----------------	--

Semi-domain-independent Strategies

sweep	Enumerate objects connected each to the next by the same relation.
follow a path	Traverse a natural connection between parts of the knowledge to be described.

Figure 2: Strategies.

that bears the same relationship to the just-mentioned one, until there are none left; for example, “to the left of the window is a stove and then a refrigerator” is a *sweep-left*. Similar constructions may be found in other domains. A description of one’s day may start at some event and mention the next event and the event after that. In this case, the relationship is temporal, rather than physical.

Metastrategies select strategies based on the *context*, which comprises:

Local Context

- What has just been said.
- What is *immediately* connected to the current item in the knowledge base.
- What strategies are applicable.

Global Context

- Discourse and speaker parameters. (For example, a speaker’s propensity to mention objects to the right before objects to the left.) It is here that anything considered a global goal would be encoded.
- The completeness criterion.

In future implementations, the context may also involve some model of the hearer. This would be part of

the local context: what one knows about one’s hearer and about what one’s hearer knows changes, particularly under the assumption that hearer and speaker are engaged in a two-way interaction.

A strategy conflict set contains whatever strategies are currently applicable. If it contains a single strategy, that one is selected. If more than one is in the set, metastrategies may resolve the conflict by selecting one or by combining some or all of them. Finally, if no strategy presents itself, the metastrategies apply a default strategy.

The metastrategy *find-“interesting”-link* is triggered after the introduction of a new topic, which has links to many other items in the knowledge base. What is “interesting” or salient depends on:

- **The domain:** In spatial description, objects that are large or have many features are interesting.
- **The structure of the domain:** Objects that are more connected to other objects are more interesting.
- **The local context:** If a window has just been mentioned, there is reason to mention other windows.
- **The global context:** There may be an inclination to mention furnishings but not structural features of the house, or vice versa; there may be differing levels of detail required.

Some metastrategies combine strategies. Such metastrategies apply when several strategies are appropriate at some point in the description and there is a felicitous way to combine them. *circular-sweep* is an example: it combines a number of sweep strategies (*sweep-left*, *sweep-right*, and *sweep-under*), and includes additional orientation strategies to orient hearers between sweeps. *kitty-corner* is a metastrategy that is used to describe a room in which the most salient feature is diagonally opposite the current location. The object that is "kitty corner" from it is mentioned first, and the rest of the room is described in relation to it.² The first fragment in figure 1 exemplifies the *kitty-corner* metastrategy.

find-new-topic is the default metastrategy just in case there is nothing "interesting" to say. In a spatial domain, the default is to select an object to describe using spatial proximity.

Example of description organization

In this section, I describe the organization of a fragment of a description in my corpus. While the system is not yet fully implemented to handle all the details, this example is sufficiently complex to show the operation of the architecture in selecting appropriate strategies and metastrategies. Though the strategies used are simple, complex choices, varying with context, are made. On the following page are the text and a sketch of the area being described.

In the fragment in figure 3, the speaker is describing the bedroom he shares with his wife Carol. Each line of the fragment, in most cases, is the result of a single strategy.

The sketch in figure 4 of the bedroom is provided as an aid to the reader in understanding John's description. There is no corresponding representation in the system.

The global context used by this speaker includes parameters that predispose him to mention rather than ignore room furnishings, as well as "stuff"—small articles than can be found in, on, or near pieces of furniture. The strategy *mention-stuff* is a particular form of *describe-object*, and as such is concerned with mentioning associated properties of the object rather than physical proximity; this is suggested by the speaker's typically using the preposition "with," rather than an obviously spatial one.

The global context also of course includes the completeness criterion which is unsatisfied throughout this stretch of text. The fragment starts when the speaker has just finished describing the kitchen and the next spatially proximate thing is the door to the bedroom. When the node to be described is a physical object,

²For a fuller treatment of how the system computes deictic and other spatial terms like "kitty corner," "left," and "right," see (Sibun & Huettner, 1989).

the *mention-object* strategy is always available; because this object is a door, *mention-room* is available to talk about the room that the door leads into. The metastrategies resolve this conflict in favor of the more particular *mention-room* {1}.

Because the last strategy used was *mention-room*, *mention-salient-object* becomes available; if there is a particularly "salient" object in a room, descriptions can then be organized around it. As it happens, salience in this domain tends to depend primarily on size; the kingsize bed fits the bill {2}. There are two objects spatially proximate to the bed; one is selected and the strategy *mention-object* is used to mention the endtable {3}. The endtable is connected to several other items in the knowledge base, but because there is a context parameter to mention "stuff," this is what happens next {4}.

Now, there are two unmentioned objects spatially proximate to the endtable—the window and the "wall" (which is actually a covered-up chimney). The "wall" is mentioned because, like furnishings, it has extent in the room, and the context disposes the process toward mentioning such objects {5}. The local context keeps track of what has just been mentioned; another feature it records is the direction in which the spatial proximity links have been followed. The "wall" is next along the trajectory from the bed through the endtable.

The "wall" is spatially linked to three things: the window on the endtable side of it, the window that is along the same trajectory that has been followed, and the chests, which are also along that trajectory. The window along the trajectory is the choice selected from these three for two reasons: there is a tendency to maintain trajectories;³ and the last thing mentioned is part of the structure of the room, as is the window, but not the chests. But this selection of the window presents a problem (at least, we can infer that it presents a problem to the speaker), because this window is connected via a similarity link to the previous window, which has not been mentioned. So the speaker performs a repair consisting of backing up, mentioning the overlooked window, reiterating mention of the wall, and, finally, mentioning the selected window {6}.

While the next window is a promising candidate because it is the same sort of thing as that just mentioned, the parameter for mentioning furniture overrides this, and the cedar chests come next in the description {7}, followed by their associated "stuff" {8}. The window is again available but so is the bureau, which is the preferred choice because it is furniture {9}. The bureau has spatial proximity links to two windows and the "small thing," as well as links to its "stuff," so the set of available strategies comprises ones that mention each of these things. A metastrategy can resolve this

³Ullmer-Ehrich (1982) notes a similar tendency in the dorm room descriptions that she collected.

and then there's the bedroom {1}
 and there's a huge kingsize bed {2}
 and there's an endtable next to it {3}
 with a lamp and a clock radio {4}
 and some stuff of mine
 like the Boston University cup
 and...then there's a wall-- {5}
 and then there's a window behind that {6}
 and there's a wall, and there's another window
 and there's some...cedar chests of Carol's {7}
 that have blankets and sheets in them {8}
 and...there's her bureau {9}
 in the middle of two windows on either side {10}
 with all of her makeup on top of it and clothes
 and there's...a small thing with all her clothes
 and there's another great big bookshelf
 and all her spare books
 and there's a small end table over on her side
 um...a small digital clock and more kleenex
 and...OK

Figure 3: John's description.

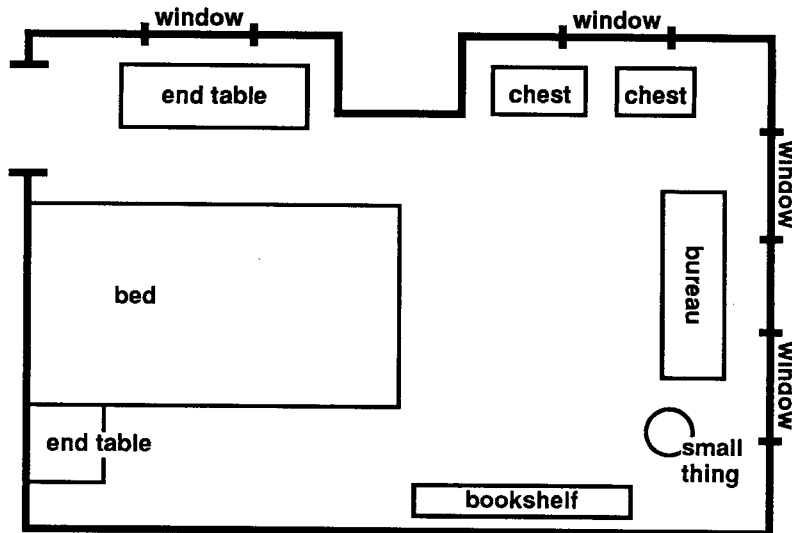


Figure 4: Sketch of bedroom.

conflict by realizing that strategies for saying that there is an object of the same sort on two different sides of the current object can be combined by saying that the bureau is between the two windows {10}.⁴

The description for the rest of the room continues in a manner similar to that already discussed.

Incremental generation

The organizer described here is composed of strategies that are often responsible for sub-clausal units of text; furthermore, the strategies have already imposed an organization on the text, obviating the need for the generator to do more than *enforce grammaticality*. The model of local text organization I am developing is coupled with a model of *incremental generation*, in which the increments are often smaller than a sentence. (Garrett's investigation of speech errors (1975) constitutes early work in this area; De Smedt & Kempen (1987), and Kempen & Hoenkamp (1987) discuss a similar, more fully-developed incremental generation project.) A typical generator produces a sentence at a time. However, spoken text is replete with restarts, fragments, and ungrammaticalities. This suggests that not only do people organize their text incrementally, but generate it in increments as well.

Usually, an incremental generator is successful in generating grammatical text. The text

then /
in the kitchen /
there is a large window

is the result of three sequentially operating strategies: *introduce/shift-attention*, *situate*, and *mention-salient-object*. An incremental generator will be able to produce this text in the increments specified by the strategies.

A system that generates in strategy-sized increments can result, in principled ways, in ungrammatical text. A common error, exemplified by

and a lot of other people /
who /
I wasn't quite sure what they did

can be explained by the operation of the strategies *mention-object*, *add-clausal-modifier*, and a set of strategies to express some additional information, which, because a dependent clause has already been started, happens to result in the ungrammatical resumptive pronoun "they."

Locally-organized, occasionally ungrammatical text may be prototypical, but it is certainly not the only sort

⁴Note that the previous window conflict was not resolved by saying that the "wall" was between the two windows. The difference can be explained by the observation that the "wall," despite its extent into the room, is a structural object, while the bureau is furniture.

of text we wish to generate. To be comprehensive, my system will require some capability to post-process text after it is organized and before it is generated (Hovy, 1989, suggests a post-processing model). Output of my system might also be appropriate input for a text revision system (e.g., Meteor & McDonald, 1986).

Related research

There are many other projects whose goal is to organize, or plan, extended text. The main difference between these and mine is flexibility and level of organization: most text planning systems rely on global structures to organize paragraph-sized text. These structures, which are usually schemas or plans, constrain the text to reflect particular rhetorical and domain relations, but at the expense of flexibility. My system builds structure locally with no recourse or reference to an overall structure.

Through analysis of written texts, McKeown (1985) has developed a small set of schemas, built from rhetorical predicates, that would provide types of descriptions (for example, constituency) for objects in a knowledge base. Paris has extended this work to use a process trace which derives its structure in part from the domain (Paris & McKeown, 1987; Paris, 1988). This alternative strategy builds and traverses a path through the knowledge base when it is determined, on the basis of a sparse model of user expertise, that a process description of an object is more appropriate than a declarative one; the two strategies may be interleaved. Dale (1989) similarly organizes text by means of a domain structure, in this case, a recipe plan.

Rhetorical Structure Theory (Mann & Thompson, 1987) is also drawn from an analysis of written texts; it differs from McKeown's work in that it is composed of a large number of rhetorical relations, rather than a small number of schemas. RST may thus be more flexible, but it is still assumed that the relations will be combined into a single global tree covering an extended text.

While most text planning systems have worked on producing single texts in response to queries, some research has been more particularly concerned with interactive text. Much recent work in this area has been for explanation systems (e.g., Maybury, 1989), and some of this work explicitly addresses allowing a human user to ask follow-up questions (Moore & Swartout, 1989). However, such systems still build and use global structures for extended texts.

An argument is sometimes made that global structure is needed to capture high-level rhetorical goals in the output text (see Appelt, 1985); Gricean Maxims (Grice, 1975) are often invoked. But, as Hovy (1988) points out, "be polite" is not a decomposable goal; the objective of being polite is achieved through local decisions. Such local decisions can comfortably be integrated with a model of local organization of text.

Acknowledgements

I thank Bonnie Webber and David Chapman for much useful discussion in the preparation of this paper. I also thank Bruce Leban for helping to make the Mac draw the picture. This work was supported in part by the Air Force Systems Command, Rome Air Development Center, Griffiss AFB, New York, 13441 under contract No. F30602-85-C-0008, as part of the Northeast Artificial Intelligence Consortium (NAIC).

References

- Appelt, D., "Planning English Referring Expressions." *Artificial Intelligence* 26, pp 1-33, 1985.
- Dale, R., "Generating Recipes: An Overview of Epicure." In *Extended Abstracts Presented at the Second European Natural Language Generation Workshop*, pp 7-13, 1989.
- Davis, R., "Meta-rules: Reasoning about Control." *Artificial Intelligence* 15, pp 179-222, 1980.
- De Smedt, K. and G. Kempen, "Incremental Sentence Production, Self-correction, and Coordination." In G. Kempen (ed), *Natural Language Generation: New Results in Artificial Intelligence, Psychology and Linguistics*, Martinus Nijhoff Publishers, pp 365-376, 1987.
- Frohlich, D. and P. Luff, "Applying the Technology of Conversation to the Technology for Conversation." In *Proceedings of the Computers and Conversation Symposium*, University of Surrey, England, 1989.
- Garrett, M., "The Analysis of Sentence Production." *Psychology of Learning and Motivation*, Volume 9, pp 133-177, 1975.
- Grice, H., "Logic and Conversation," In P. Cole and J. Morgan (eds.), *Syntax and Semantics 3: Speech Acts*, pp 41-58, 1975.
- Grosz, B. and C. Sidner, "Attention, Intentions, and the Structure of Discourse." *Computational Linguistics*, (12)3, 1986.
- Hovy, E., "Two Types of Planning in Language Generation." In *Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics*, pp 179-186, 1988.
- Hovy, E., "Unresolved Issues In Paragraph Planning." To appear in *Proceedings of the Second European Natural Language Generation Workshop*, 1989.
- Kempen, G. and E. Hoenkamp, "An Incremental Procedural Grammar for Sentence Formulation." *Cognitive Science* 11, pp 201-258, 1987.
- Laird, J., A. Newell, and P. Rosenbloom, "Soar: An Architecture for General Intelligence." *Artificial Intelligence*, 33, pp 1-64, 1987.
- Levinson, S., *Pragmatics*. Cambridge University Press, 1983.
- Linde, C., *The Linguistic Encoding of Spatial Information*. Doctoral Dissertation, Columbia University, 1974.
- Mann, W. and S. Thompson, *Rhetorical Structure Theory: A Theory of Text Organization*. Report ISI/RS-87-190, Information Sciences Institute, 1987.
- Maybury, M., "Enhancing Explanation Coherence with Rhetorical Strategies." In *Proceedings of the European Association of Computational Linguistics*, pp 168-173, 1989.
- McKeown, K., *Text Generation*. Cambridge University Press, 1985.
- Meteer (Vaughan), M. and D. McDonald, "A Model of Revision in Natural Language Generation." In *Proceedings of the 24th Annual Meeting of the Association for Computational Linguistics*, pp 90-96, 1986.
- Meteer, M., D. McDonald, S. Anderson, D. Forster, L. Gay, A. Huettner, and P. Sibun, MUMBLE-86: *Design and Implementation*. COINS Technical Report 87-87, University of Massachusetts, 1987.
- Moore, J. and W. Swartout, "A Reactive Approach to Explanation." In *Proceedings of the International Joint Conference on Text Generation IJCAI*, 1989.
- Paris, C., "Tailoring Object Descriptions to a User's Level of Expertise." *Computational Linguistics* 14(3), pp 64-78, 1988.
- Paris, C. and K. McKeown, "Discourse Strategies for Describing Complex Physical Objects." In G. Kempen (ed), *Natural Language Generation: New Results in Artificial Intelligence, Psychology and Linguistics*, Martinus Nijhoff Publishers, pp 97-115, 1987.
- Reichman, R., *Getting Computers to Talk Like You and Me*. MIT Press, 1985.
- Sibun, P. and A. Huettner, "Spatial Deixis in Generating Descriptions." COINS Technical Report 89-34, Department of Computer and Information Science, University of Massachusetts, 1989.
- Ullmer-Ehrich, V., "The Structure of Living Space Descriptions." In R. Jarvella and W. Klein, eds., *Speech, Place, and Action*, John Wiley & Sons, Ltd., 1982.