Choosing Rhetorical Relations in Instructional Texts: The Case of Effects and Guidances

Leila Kosseim
kosseim@iro.umontreal.ca

Guy Lapalme
lapalme@iro.umontreal.ca

Département d'informatique et de recherche opérationnelle
Université de Montréal
CP 6128, Succ. Centre Ville
Montréal, Québec, Canada
H3C 3J7

Abstract

In this paper, we address the problem of planning the textual organization of instructions. We take the view that natural language generation (NLG) is a mapping process of different levels of conceptual and textual representations. Within this framework, we consider the mapping between the text's semantic representation and its rhetorical structure. We argue that such a mapping is not direct, but rather many-to-many, and give concrete examples of such a phenomenon in instructional texts. We then discuss the case of two semantic elements (called here *semantic carriers*), namely effects and guidances, we determine by what rhetorical relations they are most frequently realized in instructional texts, and finally, we show how such a mapping can be performed automatically within a text generation system.

word count:

```
detex ewg.tex | wc
793 4351 28508
```

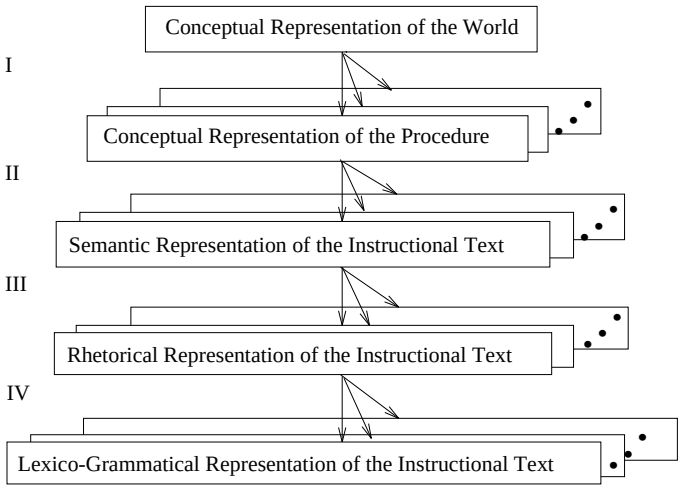

Figure 1: Mapping Levels for the Generation of Instructional Texts

Introduction

In this paper, we view the generation of instructional texts as the multilevel mapping process of figure 1[1]. Within this framework, we present a solution to the problem of selecting the most appropriate rhetorical structure to communicate a semantic representation; that is, how to perform the mapping process III of figure 1. We discuss why, in some cases a semantic representation is conveyed through a rhetorical relation, while under other conditions another relation is preferred. We show what many-to-many relation exists within a monolingual (French) context, we discuss the types of constraints that influence the choice of a rhetorical structure and propose a set of guidelines for the automatic selection of RST relations [Mann and Thompson, 1988].

1 Generating Instructional Texts

We view the generation of instructional texts as the mapping process of figure 1.

Mapping Process I: From a conceptual representation of the world (for example, a library of uninstantiated operations schemas), some planning process selects, links and instantiates knowledge into a conceptual representation of the procedure to be described. This step is generally performed by a task planner similar to NOAH [Sacerdoti, 1977] (*cf.* [Mellish, 1988, Dale, 1992, Vander Linden, 1993, Kosseim and Lapalme, 1994]). This choice follows psychological evidence that a procedure's representation is hierarchical in structure and contains instantiated schemas of operations [Dixon et al., 1988, Britton et al., 1990, Donin et al., 1992].

[1]This model assumes a linear generation process.

Mapping Process II: From the conceptual representation of the procedure, the content of the text (what will be included and what will be left unsaid) is selected. This mapping produces a semantic representation of the instructional text. This step has not received much attention in previous work (but see for example [Kosseim and Lapalme, 1994] for an attempt), probably by fear that computational linguists start from rather subjective conceptual representations of procedures.

Mapping Process III: From the semantic content of the text, the rhetorical structure is selected. In previous work (eg. [Mellish, 1988, Dale, 1992]) this mapping is mostly considered one-to-one rather than many-to-many and is not linguistically motivated. It is precisely this step that is discussed in this paper: what rhetorical relations are used to convey a semantic representation, and what contextual factors influence the choice of a preferred one.

Note that when using a constructive RST text planning technique (eg, [Moore and Paris, 1988, Hovy, 1993]) steps II and III are combined into one; the content of the text is selected through rhetorical relations, thus missing a level of mapping.

Mapping Process IV: Finally the rhetorical structure of the text is mapped onto the most appropriate lexico-grammatical representation (eg. [Vander Linden, 1993]).

To analyze the "natural" mapping between the semantic and the rhetorical levels of instructional texts, two approaches are available: studying human professional technical writers performing their work, or studying the resulting texts. Interrogating professionals at work is efficient only if they are conscious of their choices and can justify them. However, according to [Rettig, 1991], a great number of human instruction writers are not professional technical writers, but rather the technicians or engineers who developed the product to be described. These people do not generally enjoy technical writing and often do not know how to do it efficiently [Puscas, 1989]. The claims made here are therefore based on a corpus analysis of French instructional texts. In total, the corpus is composed of about 13,300 words.

2 Mapping Semantics onto Rhetorics

In [Delin et al., 1993], it has been argued that in a multilingual instructional text environment, the same information can be conveyed using different RST structures depending on the language of communication. More generally, many researchers have argued that within a monolingual environment the mapping between the semantic and the rhetorical levels is many-to-many [Moore and Pollack, 1992, Korelsky and Kittredge, 1993]. In order to generate high quality texts, it is necessary to have linguistically motivated guidelines on how to organize a text's content; that is how to map a semantic representation onto the most appropriate RST structure.

In French[2] instructional texts, sentences like the following appear quite often.

[2]The examples given throughout the paper have been translated from French in order to simplify the text. We have tried as much as possible to keep the same grammatical form in order to illustrate our points. However, we cannot over-stress that the analysis was performed on French and not on English texts.

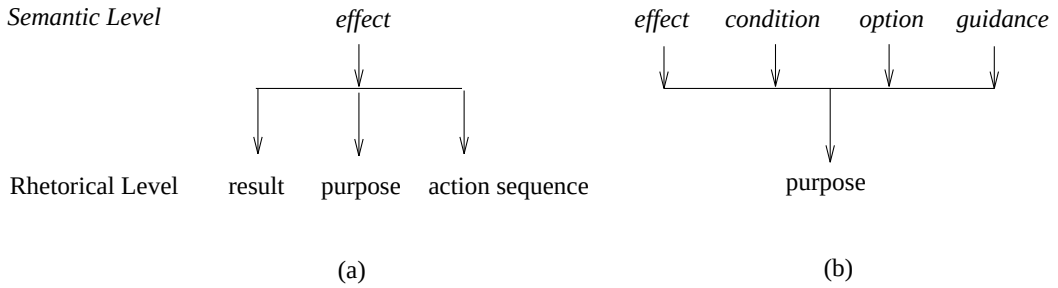

Figure 2: Many to Many Mapping

(1) a. Plug the electrical cord of the video-tape recorder in a 120V outlet and press on the POWER button. *The* POWER *light is turned on and the clock starts to blink.*

b. Screw the screw-cap on the lamp-shade holder *so that you do not lose the it.*

c. *You can see the volume level* by observing the red bar on the 15-bar scale displayed on the screen.

In these three examples, the same semantic information is conveyed by the expression in italics: it expresses the *effect* of some action. However, these effects are communicated through different RST relations. In the case of 1a, a **result** is used; in 1b a **purpose** is used; and in 1c a **temporal sequence** of actions is used. This mapping is shown in the diagram of figure 2a.

Inversely, one rhetorical relation can be used to convey different semantic informations. For example, the relation of purpose in 1b communicated the *effect* of an action; while in 2a it communicates a *condition* on an action; in 2b, it communicates the *optional nature* of an action; and finally in 2c it communicates a *guidance* of how to perform an action. This is illustrated in figure 2b.

(2) a. *For [checking] an ordinary plug,* [...] touch the copper screw with the clip of the checker.

b. Pull the wheel and the tire; *to ease the task,* firmly press on the side of the tire with your foot.

c. Turn this knob clockwise and counter-clockwise *to minimise interference.*

The question of which semantic element is conveyed by a rhetorical relation is not of our concern here but demonstrates that, within a single language, the mapping between a text's semantic representation and its rhetorical structure is many-to-many and does occur rather frequently.

3 Factors Constraining the Choice of a Rhetorical Relation

Instructional texts possess rather stereotypical semantic content and rhetorical structures. From the semantic point of view, 8 procedural *semantic carriers* are typically found: sequential, co-temporal and eventual operations, options, material conditions, guidances, effects[3] and operation prevention [Kosseim and Lapalme, 1994]. Except for guidances, these semantic carriers are "elements of meaning" that are mapped onto the satellite of an RST relation. These semantic carriers are typically conveyed through 6 rhetorical relations: temporal sequence, action concurrency, means, c-condition[4], purpose and result [Rösner and Stede, 1992, Vander Linden, 1993]. The choice of a rhetorical relation depends on several factors:

The content of the semantic representation: Obviously, the most important factor in determining what rhetorical relation to use is what semantic information we wish to convey. For example, a *condition* cannot be conveyed through a result or an action concurrency. For a particular semantic carrier a set of acceptable rhetorical relations must be determined and as we have seen, these sets are not mutually exclusive.

The structure of the semantic representation: We view the semantic representation of the text as a tree structure resembling the structure of the conceptual representation. Its structure does influence how the information is to be conveyed; for example, an operation[5] that is divided into a large number of sub-operations influences the rhetorical choice as a sub-instruction will probably be specified.

Co-occurrence constraints: In instructional texts, some rhetorical relations seem to co-occur; while some combinations are simply never found.

For example, if two equivalent conditions are to be presented, a c-condition and a result will convey this information: the easiest condition to verify is conveyed through a c-condition; while the other uses a result. As in:

(3) If they [the screws] have an "L" mark, *they have a left winding*, and you must unscrew them [...]

However, one will never find a temporal sequence of actions related by a means relation to a concurrency. This unlikely form is shown in figure 3a. To convey the same information, a purpose related to a concurrency is preferred, as in figure 3b.

A model of the reader's knowledge and intentions: What the reader believes about the operations and states of the procedure and her pursued goals greatly influence how information is conveyed in the text. For example, in:

(4) a. *If you wish a thicker line,* stay on the glass longer so that more paint can flow.

[3] In [Kosseim and Lapalme, 1994], the term *causality* was used.

[4] A *c-condition* combines RST's relations of circumstance and condition. It is what [Rösner and Stede, 1992] and [Vander Linden, 1993] call a *precondition*, but we prefer to use this term in its AI planning definition. The relation of *means* refers to [Kosseim and Lapalme, 1994]'s *enablement*.

[5] We use the term *operation* in all levels above and including the semantic representation of the text; and the term *action* for all levels below.

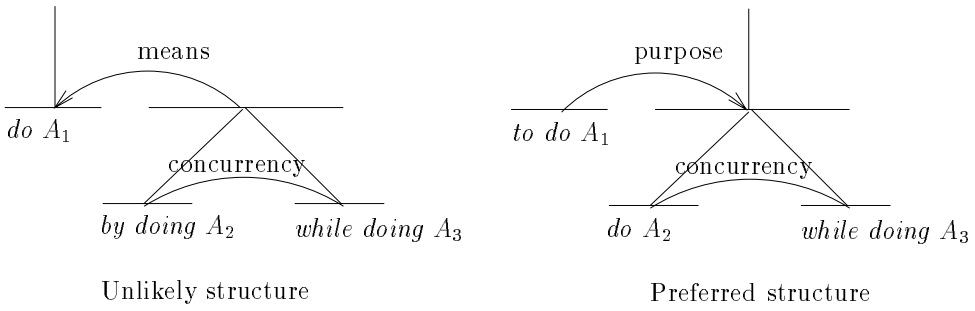

Figure 3: Preferred Rhetorical Structure

b. *To have a thicker line*, stay on the glass longer so that more paint can flow.

If the two relations in italics present the semantic carrier of option, a relation of c-condition (4a) is prefered for novice readers as the optional nature of the action *stay longer* is explicit. A relation of purpose (4b) does not convey the optionality as explicitly and can be mistakenly interpreted by a novice reader as a mandatory goal to be achieved.

The model of the reader should be allowed to be incomplete and inconsistent compared to the reality, and should be updated dynamically as the text is generated.

Specifications of the nature of the procedure: This constraint takes into account functional characteristics of the operations and states of the procedure to select a rhetorical relation. This includes the optionality and degree of desirability of an operation (if an optional line of operations is generally desirable, it will be conveyed differently than one rarely chosen), the level of danger of a negative operation, the internal/external status of states, ...

Taking these constraints into consideration, let us see how the choice of a rhetorical structure can be performed.

4 Effects and Guidances in Instructions

Let us consider now how to convey two semantic carriers often found in instructional texts: effects and guidances. We define that a linguistic expression conveys an *effect* if:

- it communicates a state brought about by a user[6] or a non-user operation (what is usually called a postcondition in AI planning) or
- it communicates a mandatory operation O_1 that is *generated* (in the sense of [Goldman, 1970]) by another operation O_2 expressed in the text and O_2 does not influence how O_1 should be performed.

[6] By *user*, we mean the main human agent of the procedure. In instructional texts, often, but not always, the user is the reader of the text.

Rhetorical Relation	**Nb. of Occurrences**	**Proportion**	**Example**
result	92	69 %	[Do *A*.] *E will be done.*
purpose	39	29 %	*To achieve E,* [do *A*].
temporal sequence	3	2 %	*Do E* [by doing *A*].
total	134	100 %	

Table 1: Mapping of Effects onto Rhetorical Relations

On the other hand, let us define a *guidance* as the information conveyed by a linguistic expression that:

- communicates a mandatory operation O_1 that is *generated* by another operation O_2 expressed in the text; but this time, O_2 does influence how O_1 should be performed.

For example, in:

(5) Turn this knob clockwise and counter-clockwise *to minimise interference.*

the expression in italics conveys a guidance, as the operation O_1 (*minimise interference*) is generated by O_2 (*turn*) but O_1 influences how O_2 should be performed. However, in:

(6) Screw the screw-cap on the lamp-shade holder *so that you do not lose the it.*

an effect is conveyed as the goal *not losing the screw-cap* does not influence how the cap is screwed on the lamp-shade holder.

4.1 Mapping Effects onto RST Structures

As illustrated in example 1 and in figure 1a, in French instructional texts effects can be communicated through three different RST relations: **results**, **purposes**, and **temporal sequences**[7]. Table 1 shows the frequency of each relation in our corpus.

Effects do not seem sensitive to the structure of the text's semantic representation and to co-occurrence constraints. The model of the reader and the nature of the procedure do however influence how effects are presented.

Results: Results are the most common choice to convey an effect (69%).

If the effect is an external reaction from the device (it is external to the device and generated by a non-reader operation) then a result is the best alternative. This is the case in example 1a. Indeed, the reader does not usually want to achieve an external effect but an internal one. The external effect is only a window on what is happening inside the device. In example 1a, the reader's goal is not to have the POWER light on

[7] in [Kosseim and Lapalme, 1994], only the first 2 relations are considered as temporal sequences are rather rare.

and the clock to start blinking; she wants the device to receive electricity in order to be functional. Using a purpose in that case would therefore be unnatural as in:

(7) *To turn on the* POWER *light and have the clock start blinking*, plug the electrical cord of the video-tape recorder in a 120V outlet and press on the POWER BUTTON.

A result is also used if the reader cannot guess that the effect is desirable. In that case, a result is prefered over a purpose because it ensures that she will interpret the actions specified in the nucleus as mandatory and not optional. Consider, for example, the sentences:

(8) a. Change the gear rapidly in each position [...], *this makes the transmission liquid circulate.*

b. Change the gear rapidly in each position [...] *to make the transmission liquid circulate.*

In 8a, the reader is not inclined to consider the degree of desirability of the effect and thus considers the *change gear* action as mandatory. However, in 8b, the reader can[8] evaluate if she wishes to attain the effect and in that case should consciously decide to *change gear.*

Purposes: Purposes are also frequently used to convey an effect (29%) (see example 1b). As seen above, purposes are not used for external effects but used when the reader knows or can guess that the effect is desirable. A purpose implicitly gives the reader the option to execute or not the line of actions specified in the nucleus. Using such a relation is therefore only adequate if the reader understands that the effect is desirable and will choose on her own to execute the next line of actions.

Temporal Sequences: Temporal sequences are sometimes used in French instructions to communicate an effect. However, in our corpus, they are only used 2% of the time (see example (1c)); it is therefore rather difficult and unsafe to develop any heuristics for choosing such a relation, furthermore, each case could very well have used a result or a purpose. This phenomenon has also been identified for English instructions by [Delin et al., 1994] in the example:

(9) a. Pull down and remove the white plastic tray that holds the video cable and *unpack the cable.*

b. Pull down and remove *to unpack the video cable.*

4.2 Mapping Guidances onto RST Structures

Guidances can be conveyed through two RST relations: **purposes**, and **means**. For example:

(10) a. *With a flat screwdriver,* scrape the dirt accumulated on the contact.

b. Adjust the belt *by pulling it by the flap.*

[8]This is not to say that she should.

Rhetorical Relation	Nb. of Occurrences	Proportion	Example
means	84	65 %	*Do G* [by doing *A*]
purpose	46	35 %	*To achieve G,* [do *A*].
total	130	100 %	

Table 2: Mapping of Guidances onto Rhetorical Relations

c. *To remove the lamp-shade holder,* press on the mantles, compress the holder and remove it from the frame.

Table 2 shows the frequency of each relation in our corpus.

Guidances are rather special semantic carriers as they do not map directly onto an RST satellite; they are relations rather than elements. A guidance can exist between two sets of operations from the semantic representation: one parent operation and its child-operations. If a relation of means presents the guidance, the satellite is made up of the child-operations and the nucleus is the parent operation. If a relation of purpose is preferred, the satellite refers to the parent operation and the nucleus to the child-operations. This phenomenon is illustrated in figure 4. In each case, a different set of operations is chosen to be in the nucleus position as the focus, or main operation, of the instruction. To explain such a choice, we have extended the notion of basic-level categories of concrete objects of [Rosch, 1978] to categories of events. Basic-level operations are used more easily than operations at other levels of the taxonomy. Basic-level operations are the most general categories for which members have the most common properties. They are therefore chosen more frequently by writers and allow readers to construct a conceptual representation of the procedure more easily. Our notion of basic-level operation follows the work of [Pollack, 1986] on domain-basic act-types, and not the work of [Goldman, 1970, Danto, 1973] on basic-level act-types.

The choice between a purpose and a means relation to express a guidance depends on the user model (basic-level operations), the nature of the procedure and co-occurrence constraints. Indeed, for co-occurrence constraints, one notices that all actions at the same abstraction level in the semantic representation are presented by the same rhetorical relation, while actions from different abstraction levels are conveyed by different relations. For example, the guidances in:

(11) To do A_{gp}, do A_p by doing A_{c1} and A_{c2}.

refer to three different abstraction levels (grandparent, parent, and child). The actions at the same level (A_{c1} and A_{c2}) are all presented by a relation of means; while the actions from different abstraction levels are conveyed by different relations.

Means: A relation of means is used most often to convey guidances (65 % of the time).

If a single child-operation is to be presented and it indicates the use of a particular instrument then a relation of means is preferred; this is the case in example 10a. In that case, the instrument is considered more important than the action of using it and

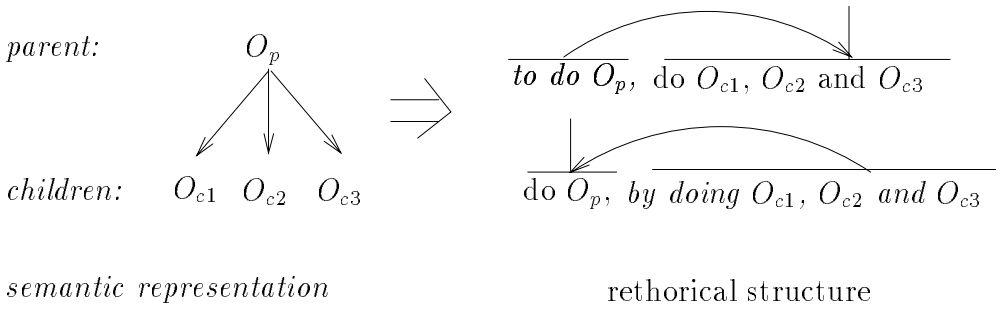

Figure 4: Guidances in Instructions

can thus be elided. Such an elision leaves no possibility to use a relation of purpose, because there will be no grammatical realization for a verb-less nucleus.

(12) ? To do A, [use] this instrument.

In the case of multiple child-operations, if at least one of the children indicates the use of an instrument, and the remaining operations are related to each other by a temporal sequence, then a means is likely to be used. For example, in:

(13) Do A *with this instrument by doing* A_2 *and doing* A_3.

A_2 and A_3 are interpreted as sequential. If they should not be, a purpose will be preferred (see below).

Finally, if a single child-operation is to be presented and the parent is a basic-level operation; the latter is generally put in focus and placed in the nucleus position. In that case, the guidance is seen top-down through a relation of means. This is the case in example 10b.

Purposes: A relation of purpose is used in all other cases (35% of the time).

If many child-operations are to be included in the text and none indicate the use of a particular instrument, then a purpose is generally preferred. Also in the case of many child-operations not related to each other by a temporal sequence but by a co-temporality, a relation of purpose will be used. This was illustrated in figure 3.

Finally, if a single child is to be presented, it does not indicate an instrument and is a basic-level operation, the relation is generally seen bottom-up by a purpose in order to put the child-operation in focus (nucleus) position.

5 Evaluation

As evaluation is becoming a necessary step in NLG research [Bates et al., 1994], we have tried, to some extent, to evaluate the above claims. We believe that a formal comparison of "natural" texts and automatically generated ones (whether through an implementation or a manual run) is a sufficient but not necessary condition to evaluate

any NLG theory. In the case at hand, to evaluate the above heuristics, we implemented them in a text generation system called SPIN[9]. This Prolog system is capable of constructing a conceptual representation of the procedure through a task planning technique similar to NOAH's [Sacerdoti, 1977]. SPIN then maps this representation onto a semantic representation of the text by applying heuristics for the selection of the text's content (see [Kosseim and Lapalme, 1994]). Finally the system maps this semantic structure onto a rhetorical structure by using presentation heuristics as the ones described above.

To evaluate the claims, we ran texts from the training corpus through the system and qualitatively commented SPIN's results. The same task was performed with texts outside the training set, including Dale's *Butter Beans Soup* [Dale, 1992] and Mellish's broken fuse example [Mellish, 1988]. The fact that the resulting semantic and rhetorical choices seem "natural" and do not include any "strange" result leads us to believe that the above guidelines are acceptable. We do realize that such an evaluation is not strict at all; but has been adopted until a better method that takes into account the richness and flexibility of natural language is developed.

6 Conclusion and Further Research

In this paper, we have taken the view that NLG is a problem of mapping between different levels of conceptual and textual representations. In this framework, we have shown, through the example of instructional texts, that the semantic representation of a text and its rhetorical structure are connected through a many-to-many mapping. We have also shown how to select the most appropriate rhetorical structure to convey a semantic representation by taking the case of effects and guidances in French instructional texts.

From the results of [Delin et al., 1993] it is clear that the choice of a rhetorical relation to convey some semantic element is dependent on the language of communication. Our research was performed on French instructions; as the examples of the paper show, the results seem applicable to English too, but we do not wish to claim such an applicability without further analysis of English instructions.

This research does not aim at finding strict and infallible rules to map the semantic representation of the text to its rhetorical structure: we developed heuristics. Indeed, through our corpus analysis, we realized that a few mappings seem to be equivalent and interchangeable within the same context. This is the case, for example, for effects conveyed through a temporal sequence of actions and through purposes or results (*cf.* section 4.1). In these cases, the different rhetorical structures seem to be due to human personal preferences that are difficult to justify and perhaps should not be.

Acknowledgments

Many thanks to all members of the SCRIPTUM research team, and especially to Massimo Fasciano for proofreading the paper. This work was supported by a NSERC scholarship.

[9]Système de Planification d'INstructions

References

[Bates et al., 1994] Bates, L., Hovy, E., and Senneff, S. (1994). (invited panelists). Panel discussion: Evaluating Generation Systems, 7^{th} International Workshop on Natural Language Generation. unpublished.

[Britton et al., 1990] Britton, B., Van Dusen, L., Glynn, S., and Hemphill, D. (1990). The Impact of Inferences on Instructional Text. *The Psychology of Learning and Motivation*, 25:53–70.

[Dale, 1992] Dale, R. (1992). *Generating Referring Expressions: Constructing Descriptions in a Domain of Objects and Processes*. ACL–MIT Press Series in Natural Language Processing. The MIT Press.

[Danto, 1973] Danto, A. (1973). *Analytical Philosophy of Action*. Cambridge at the University Press.

[Delin et al., 1994] Delin, J., Hartley, A., Paris, C., Scott, D., and Vander Linden, K. (1994). Expressing Procedural Relationships in Multilingual Instructions. In *Proceedings of the 7^{th} International Workshop on Natural Language Generation*, pages 61–70, Kennebunkport, Maine.

[Delin et al., 1993] Delin, J., Scott, D., and Hartley, T. (1993). Knowledge, Intention, Rhetoric: Levels of Variation in Multilingual Instructions. In Rambow, O., editor, *Proceedings of the ACL Workshop on Intentionality and Structure in Discourse Relations*, pages 7–10, Ohio State University.

[Dixon et al., 1988] Dixon, P., Faries, J., and Gabrys, G. (1988). The Role of Explicit Action Statements in Understanding and Using Written Directions. *Journal of Memory and Language*.

[Donin et al., 1992] Donin, J., Bracewell, R., Frederiksen, C., and Dillinger, M. (1992). Students' Strategies for Writing Instructions, Organizing Conceptual Information in Text. *Written Communication*, 9(2):209–236.

[Goldman, 1970] Goldman, A. (1970). *A Theory of Human Action*. Princeton University Press, Princeton, NJ.

[Hovy, 1993] Hovy, E. H. (1993). Automated discourse generation using discourse structure relations. *Artificial Intelligence*, 63(1–2):341–385.

[Korelsky and Kittredge, 1993] Korelsky, T. and Kittredge, R. (1993). Towards stratification of RST. In Rambow, O., editor, *Proceedings of the ACL Workshop on Intentionality and Structure in Discourse Relations*, pages 52–55, Ohio State University.

[Kosseim and Lapalme, 1994] Kosseim, L. and Lapalme, G. (1994). Content and Rhetorical Status Selection in Instructional Texts. In *Proceedings of the 7^{th} International Workshop on Natural Language Generation*, pages 53–60, Kennebunkport, Maine.

[Mann and Thompson, 1988] Mann, W. and Thompson, S. (1988). Rhetorical Structure Theory: towards a functional theory of text organization. *TEXT*, 8(3):243–281.

[Mellish, 1988] Mellish, C. (1988). Natural Language Generation from Plans. In Zock, M. and Sabah, G., editors, *Advances in Natural Language Generation*, Communi-

cations in Artificial Intelligence Series, chapter 7, pages 131–145. Pinter Publishers, London.

[Moore and Paris, 1988] Moore, J. and Paris, C. (1988). Constructing Coherent Text using Rhetorical Relations. In *Proceedings of the* 10^{th} *Annual Conference of The Cognitive Science Society*, pages 637–643, Montréal.

[Moore and Pollack, 1992] Moore, J. and Pollack, M. (1992). A Problem for RST: The Need for Multi-Level Discourse Analysis. *Computational Linguistics*, 18(4):537–544.

[Pollack, 1986] Pollack, M. (1986). *Inferring Domain Plans in Question-Answering.* PhD thesis, University of Pennsylvania. SRI technical Report SRIN-403.

[Puscas, 1989] Puscas, M. (1989). A Survey of Technical Computer Users Resulting in Guidelines for the Development of Technical Computer Documentation. In *Conference Proceedings*, pages 49–65, Baltimore, MD. SIGDOC, The Association for Computing Machinery.

[Rettig, 1991] Rettig, M. (1991). Nobody Reads Documentation. *Communication of the ACM*, 34(7):19–24.

[Rosch, 1978] Rosch, E. (1978). Principles of Categorization. In Rosch, E. and Lloyd, B., editors, *Cognition and Categorization*, pages 27–48. Lawrence Erlbaum, Hillsdale, NJ.

[Rösner and Stede, 1992] Rösner, D. and Stede, M. (1992). Customizing RST for the Automatic Production of Technical Manuals. In Dale, R., Hovy, E., Rösner, D., and Stock, O., editors, *Aspects of Automated Natural Language Generation*, Lecture Notes in Artificial Intelligence, pages 199–214. Springler–Verlag, Trento, Italy.

[Sacerdoti, 1977] Sacerdoti, E. D. (1977). *A Structure for Plans and Behavior.* Elsevier, New York.

[Vander Linden, 1993] Vander Linden, K. (1993). *Speaking of Actions: Choosing Rhetorical Status and Grammatical Form in Instructional Text Generation.* PhD thesis, University of Colorado.