Avoiding Mis-communication in Concept Explanations

Ingrid Zukerman
Department of Computer Science
Monash University
Clayton, VICTORIA 3168, AUSTRALIA

Abstract

This paper offers a mechanism for the generation of expressions that convey a concept to a hearer. These expressions often include rhetorical devices, such as Descriptions and Instantiations. Our mechanism is based on the representation of the preconditions for the understanding of a concept in terms of failures in communication. We distinguish between two types of communication failures based on their cause: failure due to a hearer's inability to evoke an intended concept, and failure due to the hearer's lack of expertise with respect to the aspects of a concept which are necessary for the comprehension of the discourse. This categorization supports the principled selection of rhetorical devices which are tailored to the prevention of undesirable situations that result in the failure of a communicative goal. These ideas are being implemented in a discourse planner in the domain of high-school algebra.

Introduction

In knowledge acquisition settings, one agent transfers information which is believed by him/her to another agent, with the intent that the second agent learn this information. In these settings, the speaker's communicative goal is *KNOW(Hearer,item)*, where the meaning of the goal *KNOW* is that the hearer learn prescribed aspects of the item in question. This item may be a single concept, such as a procedure or an object, or a message composed of predicates and concepts. For example, the message [Distributive-Law apply-to Algebraic-Terms] is composed of the predicate apply-to and the concepts Distributive-Law and Algebraic-Terms. In an introductory algebra class, when the teacher demands that the students know this message, s/he usually means that they should know how to apply Distributive Law to Algebraic Terms. In a different setting, the requirements posed by the goal *KNOW* may differ. Educational settings are representative of situations where the item being conveyed is a single concept, e.g., *KNOW(Hearer,Conceptual-Dependency)*, since in these settings, a lesson often revolves around teaching different aspects of a concept, such as its representation and its semantics.

A necessary but not sufficient precondition for the satisfaction of the goal *KNOW(Hearer, message)* is that the subgoals *RETRIEVE(Hearer,concept)* and *KNOW(Hearer,concept)* be satisfied with respect to the concepts in the message. That is, the hearer should be able to recall the concepts intended by the speaker, and s/he should also have a level of expertise with respect to these concepts which is sufficient to understand the message. For instance, if in the above message, the speaker uses the term 'Algebraic Terms' to refer to the concept Algebraic-Terms, this term should enable the hearer to reach the intended concept in his/her memory. In addition, the hearer should know enough about Algebraic Terms to understand how Distributive Law can be applied to them. A necessary and sufficient condition for the fulfillment of the goal *KNOW(Hearer,concept)* stipulates that the subgoal *KNOW (Hearer,message)* be satisfied with respect to a set of messages the speaker considers sufficient to define the concept.

Following [Webber 1983], we distinguish between two different aspects of the goal *RETRIEVE* according to its place in the discourse: *EVOKE* – where a concept is retrieved for the first time, and *ACCESS* – which pertains to subsequent retrievals. However, we depart from Webber's *internal* perspective to discourse, which deals with relations among discourse entities, and adopt Appelt's *external* perspective, which pertains to relations among the actual objects referred to by the discourse entities [Appelt 1986].

In this paper, we specify conditions for the satisfaction of the subgoals *KNOW* and *EVOKE* with respect to concepts, and present a mechanism which generates expressions that satisfy these subgoals. In this context, we address the generation of a class of *Rhetorical Devices (RDs)* which includes Descriptions, Instantiations and Analogies. Our mechanism follows the Impairment Invalidation paradigm [Zukerman 1990a, Zukerman 1990b] which postulates that RDs are generated in order to preclude anticipated impairments to a hearer's comprehension process. Thus, the precon-

ditions for the attainment of the above subgoals are cast in terms of failures of the communication process. This approach generalizes the principle of *least collaborative effort* [Clark & Schaefer 1989] to the generation rhetorical devices. According to this principle, a speaker corrects his/her own mistakes in an effort to minimize the anticipated effort the hearer will make to understand the discourse.

In the following sections, we present the procedures for planning expressions which satisfy the subgoals *EVOKE* and *KNOW*, we examine the impairments to the satisfaction of these subgoals, and discuss the RDs which invalidate these impairments. Next, we present a simple algorithm which combines these procedures. Finally, we compare our approach with the traditional planning approach.

Planning Expressions for the Subgoal EVOKE

People use different words to refer to the same thing, and they also use the same word to refer to different things. Thus, for a particular hearer, some wordings may be more felicitous than others in achieving successful communication. Further, sometimes it may be best not to refer to an item by its name, rather, indirect means, such as similes or descriptions, may be preferable. The mechanism presented in this section takes these observations into consideration for the generation of evocative expressions.

An evocative expression is composed of two parts: (1) A *Simple Referring Expression (SRE)*, which may be a lexical item or *nil*; and (2) An *Indicative RD* which complements the SRE if it is insufficient to evoke the intended concept. Thus, an evocative expression may comprise an SRE by itself, e.g., "Look at *Mary*," an Indicative RD by itself, e.g., "Look at *the girl with red hair*," or both "Look at *Mary, the girl with red hair*." The following procedure proposes evocative expressions which combine these parts.

1. Consult the speaker's knowledge base and a model of the hearer's beliefs to propose a list of candidate SREs for a given concept (only correct SREs are extracted from the model of the hearer's beliefs).
2. For each candidate SRE, consult a model of the hearer's beliefs to ascertain whether the SRE is likely to cause an impairment to the evocation of the intended concept. If so, propose an Indicative RD that invalidates the impairment. If there is more than one suitable Indicative RD, a set of candidate evocative expressions is built by combining each Indicative RD with the SRE.

For instance, in the above example, the candidate SREs are: "Mary" and *nil*. Now, if the speaker is uncertain that the hearer knows that the name of the referent is Mary, the speaker will propose an Indicative RD, such as a partial Description, which includes sufficient attributes of Mary to enable the hearer to identify her. A null candidate SRE implies that the speaker does not know or does not wish to use the name of a concept, and thus forces the generation of an Indicative RD.

Impairments to the Evocation of a Concept

We have identified two impairments which preclude the evocation of a concept: *Lack of Connection* and *Misunderstanding*.

Lack of Connection occurs when the hearer cannot link the expression used by the speaker with the concept intended by the speaker. We anticipate Lack of Connection when one of the following conditions holds: (1) a null SRE is given by the speaker, (2) the SRE used by the speaker to refer to an intended concept does not exist in the user model; (3) the SRE exists in the user model, but it is not connected to a concept; (4) the SRE is weakly connected to the intended concept (and no other concept); or (5) the SRE is connected to the intended concept (and no other concept), but the SRE is not *accessible* in the user model, i.e., the SRE is not sufficient to activate the corresponding term in the user model[1]. The following dialogue illustrates a situation were the last condition holds:

S1: I saw Mary yesterday.
S2: [*blank stare*]
S1: You know, the one with red hair and glasses.
S2: Oh! *That* Mary!

Misunderstanding occurs when the concept retrieved by the hearer is not the one intended by the speaker. This impairment is predicted if one of the following conditions holds: (1) the SRE mentioned by a speaker is strongly connected (either correctly or incorrectly) to concepts which are different from the intended one, and either (1.1) the intended concept does not exist in the user model, or (1.2) it is not connected to the SRE, or (1.3) it is weakly connected to the SRE; or (2) the SRE is connected to several concepts, including the intended one, and either more than one concept is primed, i.e., the proposed SRE is ambiguous, or the 'wrong' concept is primed. For instance, in the above dialogue, the first condition would hold if the hearer and the speaker did not know the same Mary.

Invalidating Evocation-related Impairments

The invalidation of an evocation-related impairment requires an Indicative RD which identifies the intended concept uniquely. This RD is composed of a set of concepts which are related to the intended concept. The level of detail of an Indicative RD must be sufficient to satisfy the subgoal *EVOKE* with respect to the concepts mentioned in it, i.e., the fulfillment of the subgoal *KNOW* is not necessary with respect to these concepts.

[1]This condition is stated for the sake of completeness, but has not been implemented.

We distinguish between two types of Indicative RDs: *positive* and *negative complement*. An Indicative RD is positive, if the selected concepts indicate what the intended concept *is*, as opposed to what it *is not*, e.g., "Knowledge Representation, *where we studied Predicate Calculus and Semantic Nets.*" It is negative complement, if it contains the complement of the intended concept, i.e., all the concepts that the intended concept is not, e.g., "Knowledge Representation, *which is not Search or Automated Reasoning.*"

In order to generate a positive Indicative RD for an intended concept, we must present a set of concepts which identify uniquely this concept, i.e., we must find a set of concepts, S, such that the intersection of the concepts outside S to which all the concepts in S are connected consists only of the intended concept and no other concept. According to Grice's Maxim of Quantity [Grice 1975], S should be the smallest set which satisfies this constraint. However, this requirement makes our problem NP-complete (provable by reduction to the graph covering problem). In this research, we apply a heuristic similar to the one used in [Appelt 1982, Dale 1990], whereby a uniquely identifying set of concepts is built by iteratively adding concepts to an initial identifying concept, until the intended concept is the only one connected to all the concepts in the resulting set of concepts.

Negative complement RDs may be generated only if the following conditions hold: (1) the hearer is aware of the existence of the intended concept, and (2) the situation at hand is a 'closed world' situation, where the number of candidate concepts to which an SRE could be referring is small. For example, if the only topics covered in an Artificial Intelligence course were Knowledge Representation, Search and Automated Reasoning, a mediocre student may believe that all three topics are possible meanings of the SRE 'Knowledge Representation'. In this case, a lecturer may identify the intended concept by saying "Knowledge Representation, *which is not Search or Automated Reasoning,*" provided the student is unlikely to experience evocation-related impairments with respect to the terms in the Indicative RD.

In Lack of Connection, the hearer has no link between the given SRE and other unintended concepts. Hence, this impairment can be invalidated only by means of a positive Indicative RD. Misunderstanding may be invalidated either by means of a positive Indicative RD or a negative complement Indicative RD. Both types of RDs single out the intended concept. However, negative complement RDs also address directly a hearer's beliefs which are related to the intended concept and the SRE, while positive RDs address these beliefs only indirectly. For instance, in the Artificial Intelligence example, the negation of the erroneous meanings will both single out the intended concept and contradict the erroneous beliefs. On the other hand, a perspective, such as "Knowledge Representation, *where we studied Predicate Calculus and Semantic Nets,*" does not directly contradict the erroneous beliefs. Note that if all the possible meanings of a lexical item are correct, a negative complement RD should be limited to a negation of intent, e.g., "He was holding a crook, *not the criminal kind.*" However, if some of the possible meanings are correct and some are not, the generated text must address these cases explicitly in order to avoid erroneous presuppositions. For example, consider a situation where the hearer correctly believes that 'crook' means both shepherd's staff and criminal, and s/he incorrectly believes that it also means cooking utensil. In this case, if the speaker means 'crook' in the sense of shepherd's staff, s/he should say something like "*I don't mean a criminal, and a crook is not a cooking utensil.*"

At this stage of the generation process, all the Indicative RDs which single out an intended concept are considered possible alternatives.

Planning Expressions for the Subgoal KNOW

The expertise required from the hearer for the goal *KNOW(concept)* to be satisfied is determined by the speaker's communicative intent, which in turn is influenced by the following factors: (1) the relationship between the concept and the message; (2) the relevance of the message to a main idea the speaker wishes to convey; (3) the social setting, e.g., classroom or casual conversation; and (4) the speaker's perception of the hearer's conversational goals. The speaker's communicative goal affects the expertise requirements along two dimensions: *Content* and *Strength*.

Content determines the aspects of a concept which must be known by the hearer to satisfy the goal *KNOW*. For instance, the aspects of a procedure are: sequence of steps, goal, type, objects to which it applies and conditions for its application. The aspects of an object are: components, attributes, function, internal process, type and subclasses. However, depending on the speaker's communicative goal, only some of these aspects may be required to satisfy the subgoal *KNOW*. Consider, for example, the following dialogue:

S1: I cleaned up in 500.
S2: What is 500?
S1: It is a card game.

In this dialogue, when S1 is questioned by S2 about the concept 500, he does not provide an explanation which conveys a full understanding of this concept. Rather, S1 presents only one aspect of the concept, namely its type, since he believes that this information is sufficient to enable S2 to understand the message.

Strength determines the degree of expertise the user should have with respect to the aspects selected by Content. For example, a low or medium degree of expertise would usually be required for a concept which is mentioned in passing.

Like for the subgoal *EVOKE*, we consult a model of the hearer's beliefs in order to ascertain whether an expertise-related impairment is anticipated in relation to the concept in question. If so, a *Creative RD* is proposed to prevent this situation. Each Creative RD which can perform this task is maintained as a candidate until the final selection process. For instance, an Analogy, a Description and an Instantiation are all possible candidates to explain an unknown concept. Note that the content of Creative and Indicative RDs may overlap, since a Creative RD which describes a concept is likely to identify it uniquely as well.

Expertise-related Impairments

We have identified two impairments which preclude the goal of knowing a concept: *Lack of Understanding* and *Insufficient Understanding*.

Lack of Understanding is anticipated when there does not exist in the user model a concept which corresponds to an intended concept. It entails an evocation-related impairment, since an SRE cannot point to a concept which is unknown by the hearer.

Insufficient Understanding is predicted when there exists a concept in a hearer's memory which corresponds to an intended concept, but the hearer has a low level of expertise with respect to some aspects which are required by the speaker for the comprehension of this concept. For example, the hearer is unaware of a step of a procedure or a part of an object. This impairment may occur in conjunction with a evocation-related impairment or by itself.

Invalidating Expertise-related Impairments

The invalidation of Lack of Understanding requires the generation of a Creative RD which satisfies the subgoal *KNOW(concept)*. The information included in this RD depends on the type of the concept, e.g., procedure or object, the aspects of the concept considered relevant by the speaker, the user's ability and his/her expertise. For example, if the concept in question is a procedure, such as Distributive Law, the speaker may be interested only in the sequence of its steps and in the objects to which it applies, requiring a Creative RD which conveys these aspects, such as a Description or an Instantiation. However, different users are likely to draw different inferences from the presented information, depending on their ability and expertise. These inferences must be taken into consideration when selecting an RD. For example, if the selected Creative RD is an Instantiation which shows how Distributive Law applies to an Algebraic Term, such as $2(x+y)$, a talented student will probably make all the correct inferences regarding the other objects to which Distributive Law applies, and will also generalize the procedure from the Instantiation. On the other hand, a mediocre student cannot be relied upon to make these inferences. Furthermore, s/he may draw some erroneous inferences which have to be addressed. Thus, for a mediocre student, all the relevant information must be spelled out. Finally, the subgoals *KNOW* and *EVOKE* have to be satisfied with respect to the propositions and concepts presented in a selected RD. This may require additional Creative and/or Indicative RDs. For instance, if the hearer is not proficient with respect to the multiplication of Algebraic Terms, a Creative RD must be generated to satisfy the subgoal *KNOW(*multiplication apply-to Algebraic-Terms*)* at a level of expertise considered adequate to understand the application of Distributive Law to Algebraic Terms.

Like for Lack of Understanding, a Creative RD is required to invalidate Insufficient Understanding. However, in this case, the lack of expertise may be localized to some weak, missing or erroneous constituents. Thus, the RD may be focused on these constituents. However, in order to prevent wrong implicatures, structural constraints must be applied. These constraints stipulate that if the hearer is not proficient with respect to some information items in a particular aspect, and the system decides to mention these items, then all the information items which pertain to this aspect must be mentioned. For instance, if the speaker is interested in the objects to which Distributive Law applies, and the hearer knows only that Distributive Law applies to Algebraic Terms, simply saying "Distributive Law applies to Numbers" may carry the wrong implicature that Numbers is *all* that Distributive Law applies to. This implicature is prevented by repeating what the hearer already knows, e.g., "*In addition to Algebraic Terms*, Distributive Law applies to Numbers."

The process of generating Creative RDs is implemented by presenting the information corresponding to each aspect in a list of aspects which must be known by the hearer in order to satisfy the subgoal *KNOW*. At present, this list of aspects is hand coded for each concept in a message, but at a later stage, it may be inferred from the predicates in the message. For instance, if the intended concept is a chair, and the aspects type, function and components are specified, the following text may be generated: "A chair *is a piece of furniture* (type) *which is used for sitting* (function), *and has a seat, a back and usually four legs* (components)." The main distinction between this list of aspects and schemas [McKeown 1985] or process traces [Paris 1988] is that schemas and process traces list the types of information to be presented in order to convey a concept, whereas our list of aspects constitutes desiderata which the hearer should know in order to understand a concept. These desiderata may be fulfilled in a straightforward manner by means of a Description, as in the above example, or by indirect means, such as Analogies or Instantiations. Partial descriptions are generated by omitting from a proposed Creative RD relations which, according to our model, are known by a hearer, subject to the above structural constraints. In addition, if based on a hearer's profile, a concept is considered

too abstract or difficult for the hearer, an Instantiation may be proposed.

At present, the only Creative RDs which have been implemented are Descriptions, where the constituents of a concept are also described if this is required by the subgoal *KNOW*. In order to avoid indefinite recursion, this process is repeated only once, for the constituents of the main concept. The implementation of indirect Creative RDs, such as Analogies, Similes and stand-alone Instantiations, requires rules of inference which stipulate the conditions for inferring the desiderata from these RDs.

The Main Algorithm

The main algorithm combines the procedures described in the previous sections to generate an expression which accomplishes the goals *KNOW* and *EVOKE* with respect to the concepts in a message. The algorithm receives as input a concept C and performs the following steps.

1. Create a set called *CRDs* which includes Creative RDs that achieve the goal $KNOW(C)$. If no expertise-related impairments are anticipated with respect to C, this set is empty.
2. Create a set of pairs called *SRE-IRDs*, where each pair contains an SRE and an Indicative RD that complements the SRE to achieve the goal $EVOKE(C)$. If no evocation-related impairments are anticipated between a particular SRE and C, this SRE will appear in *SRE-IRDs* with a nil RD.
3. Create a new set of RDs called *JOINRDs*, in which each element in *CRDs* is combined with each element in *SRE-IRDs*.
4. For each element in *JOINRDs*, check whether other beliefs maintained by the hearer are adversely affected by this element. If so, add to this element an RD which prevents this effect [Zukerman 1990a].
5. Select the most succinct element in *JOINRDs*.

In this procedure, expertise-related impairments are checked first, followed by evocation-related impairments. If Lack of Understanding is anticipated in the first step and Lack of Connection is anticipated in the second step, no Indicative RD is necessary, since the Creative RD proposed to invalidate the Lack of Understanding is sufficient to identify uniquely the intended concept. However, other combinations of evocation- and expertise-related impairments may require both Indicative and Creative RDs.

Let us illustrate the workings of our procedure by means of the following example. Given the message [Distributive-Law has-goal Bracket-Elimination], our procedure is activated with respect to the concepts Distributive-Law and Bracket-Elimination. Now, let us focus on Distributive Law. We assume that the only aspect of Distributive Law which is necessary to understand the message is the sequence of its operators. Hence, if this sequence is not well known by the hearer, Insufficient Understanding is anticipated, and a Creative RD containing this sequence must be presented. This may be an Instantiation, a Description, or a combination of both. Further, if an individual step is not well known, then a Creative RD must be generated for it as well. Next, candidate SREs are proposed, and evocation-related impairments are considered. In this case, the SREs are *nil* and 'Distributive Law'. The null SRE triggers Lack of Connection, calling for a positive Indicative RD. However, since the sequence of steps of Distributive Law is sufficient to identify it, no additional Indicative RD is necessary. A text such as the following may be generated if a null SRE is used: "We apply *the following procedure* to eliminate brackets." The SRE 'Distributive Law' results in the anticipation of Misunderstanding, if the hearer is conjectured to believe that the name 'Distributive Law' points to another concept, say Bracket Simplification. In this case, like for Lack of Connection, no positive Indicative RDs need to be proposed, but a negative complement Indicative RD is still an option, e.g., "*which is not Bracket Simplification.*" After the third step in our algorithm, the RDs in *JOINRD* include: (1) all the proposed Creative RDs, where each RD is accompanied by a null SRE, (2) all the proposed Creative RDs, where each RD is accompanied by the SRE 'Distributive Law', and (3) all the proposed Creative RDs, where each RD is accompanied by the SRE 'Distributive Law' and by a negative complement Indicative RD.

If a candidate SRE is a lexical item and an RD was proposed for it, the inferences from the presentation of this information may cause adverse effects with respect to a hearer's beliefs which are related to the intended concept and the SRE. These effects take place if one of the following conditions holds *and the inferences conflict with this condition*: (1) the intended concept is connected (either correctly or incorrectly) to another lexical item, i.e., it is identified with another name; or (2) the lexical item is connected (either correctly or incorrectly) to another concept, i.e., there is more than one concept associated with the lexical item. The anticipation of such comprehension impairments calls for the generation of a Revision of the information in question if it is correct, and a Contradiction otherwise [Zukerman 1990a]. Note that these conditions do not conflict with the understanding and retrieval of the intended concept, rather, they affect related beliefs. For example, if the speaker says "She was holding a crook, *which is a shepherd's staff*," a hearer who is not proficient in English may experience confusion with respect to his/her previous belief that a crook is a criminal, since a possible implicature from this statement is that a crook is *only* a shepherd's staff. Indeed, when this example was given to a non-native English speaker, he voiced this concern. The generation of a negative complement Indicative RD to invalidate Misunderstanding prevents the second of the above conditions. However,

as stated above, these RDs may be generated only if the hearer is aware of the intended meaning of the SRE.

Our algorithm terminates by selecting the most succinct of the proposed expressions. At present, we are considering a measure of succinctness based on the probability of the components of a message.

Planning versus Failure Prevention

The plan-based approach consists of applying goal-based reasoning to generate utterances that satisfy increasingly refined communicative goals [e.g., Appelt 1982, Hovy 1988, Moore & Swartout 1989, Cawsey 1990]. The failure-prevention approach, on the other hand, proposes remedial actions to prevent or overcome anticipated failures in the communication process. Its explicit representation of communication failures supports the use of a uniform mechanism to perform tasks that require special purpose mechanisms in the plan-based approach. For example, the problem of satisfying the subgoal *KNOW* may be cast both as a planning problem and as a failure-prevention problem. For this subgoal, both approaches are equivalent in the sense that they address the problem of planning discourse by trying to satisfy the preconditions for the attainment of a communicative goal. However, while in the plan-based approach the presentation of background information is addressed by adding domain-dependent preconditions to the plan operators [Cawsey 1990], in the failure-prevention approach, this task is handled through the recognition and subsequent invalidation of expertise-related impairments. Further, the generation of discourse to satisfy the subgoal *EVOKE* is an associative task which is not amenable to handling by planning. Thus, special purpose mechanisms are invoked by a planner to perform this task [Dale 1990]. In the failure-prevention approach, this task is performed by means of the recognition and subsequent invalidation of evocation-related impairments. Finally, the impairment invalidation paradigm also accounts for the generation of additional RDs that address impairments to the comprehension process due to inferences which affect other beliefs maintained by the hearer.

Conclusion

The discourse planning procedure offered in this paper generates expressions which include rhetorical devices that address comprehension problems with respect to concepts. To this effect, our mechanism relies on a characterization of preconditions to the comprehension of a concept in terms of possible failures in communication. These failures, in turn, are expressed in terms of conditions which pertain to a hearer's beliefs.

The proposed mechanism requires an accurate model of certain aspects of a user's beliefs. Thus, it is particularly suitable for highly interactive situations, and for situations where a large amount of information is available about common user beliefs and misconceptions, such as educational settings.

Our mechanism was used successfully to analyze texts in a variety of domains ranging from educational textbooks to everyday dialogues. At present, steps 1 and 4 of the main procedure have been implemented with respect to a small subset of high-school algebra, and step 2 is currently being implemented. The output of our mechanism is in the form of predicates.

References

Appelt, D.E. (1982). Planning Natural Language Utterances to Satisfy Multiple Goals, Technical Note 259, SRI International, March 1982.

Appelt, D.E. (1986). Towards a Plan-based Theory of Referring Actions. Kempen, G. (Ed.), *Natural Language Generation: New Results in Artificial Intelligence, Psychology and Linguistics*, Kluwer Academic Publishers, 63-70.

Cawsey, A. (1990). Generating Explanatory Discourse. Dale, R.; Mellish, C.; and Zock, M. (Eds.), *Current Research in Natural Language Generation*, Academic Press, 75-102.

Clark, H.H. and Schaefer, E.F. (1989). Contributing to Discourse. *Cognitive Science* 13(2):259-294.

Dale, R. (1990). Generating Recipes: An Overview of Epicure. Dale, R.; Mellish, C.; and Zock, M. (Eds.), *Current Research in Natural Language Generation*, Academic Press, 229-255.

Grice, H.P. (1975). Logic and Conversation. Cole, P.J. and Morgan, J.L. (Eds.), *Syntax and Semantics, Volume 3: Speech Acts*, Academic Press.

Hovy, E.H. (1988). Planning Coherent Multisentential Text. In Proceedings of the Twenty-Sixth Annual Meeting of the Association for Computational Linguistics, 163-169.

McKeown, K.R. (1985). Discourse Strategies for Generating Natural Language Text. *Artificial Intelligence* 27:1-41.

Moore, J.D. and Swartout, W.R. (1989). A Reactive Approach to Explanation. In Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, 1504-1510.

Paris, C.L. (1988). Tailoring Object Descriptions to a User's Level of Expertise. *Computational Linguistics* 14(3):64-78.

Webber, B.L. (1983). So What Can We Talk About Now? Brady, M. and Berwick, R.C. (Eds.), *Computational Models of Discourse*, MIT Press.

Zukerman, I. (1990a). A Predictive Approach for the Generation of Rhetorical Devices. *Computational Intelligence*, 6(1):25-40.

Zukerman, I. (1990b). Towards a Failure Driven Mechanism for Discourse Planning: a Characterization of Learning Impairments. In Proceedings of the Twelfth Conference of the Cognitive Science Society, 495-502.