Generating Peripheral Rhetorical Devices by Consulting a User Model

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

Abstract

In the process of generating text, competent speakers take into consideration the effect their utterances are likely to have on their listeners. In particular, they try to anticipate and prevent possible comprehension problems. In this paper, we propose a mechanism which emulates this behaviour for the generation of Peripheral Rhetorical Devices, such as Motivations, Contradictions and Revisions. Our mechanism anticipates the effect of a given message on a model of a listener's beliefs, and proposes these rhetorical devices to preclude possible adverse effects. As a testbed for these ideas, a discourse planner called WISHFUL is being implemented in the domain of high-school algebra.

Introduction

In the process of generating text, competent speakers/writers try to generate utterances which are best suited to attain their communicative goals with respect to a particular audience [Hovy 1987]. In particular, speakers often anticipate the effect an intended message is likely to have on their listeners/readers[1], and if this effect is adverse in some respect, the intended message is either withdrawn or complemented with rhetorical devices which counteract the effect. Although it is often the case that speakers use stereotypical information about their audience to accomplish this task, this information is usually sufficient to allow them to generate competent discourse.

Computational models for discourse generation presented to date fall into a broad spectrum, where at one end we find the plan-based approach [Appelt 1982], and at the other the schema-based approach [McKeown 1985, Kukich 1983, Paris 1988]. In the plan-based approach, a text plan which achieves a communicative goal is produced by reasoning about the effect of various speech acts on a listener's beliefs. In the schema-based approach, standard patterns of discourse are encoded using rhetorical predicates. Hybrid approaches which combine both strategies have been proposed by [Hovy 1988] and [Moore and Swartout 1989]. In particular, Moore and Swartout expand the capabilities of a discourse planner to enable it to generate clarifying explanations in a dialogue. However, both types of discourse generation strategies are based on the implicit assumption that once an information item is understood, the only modifications that will take place in a listener's beliefs result directly from this item, and hence the listener will immediately acquire the presented information. That is, these strategies do not take into consideration the effect of indirect inferences drawn from the presented information on a listener's beliefs or the indirect effect of his/her existing beliefs on the acquisition of this information. Therefore, they fail to account for rhetorical devices which address beliefs that are indirectly related to a speaker's communicative goal. For instance, in the sample text in Figure 1 (shortened version from [Lynch et al. 1979]) the communicative goal is for the listener to know the distributive law, however, the discourse which accomplishes this goal is preceded by a Contradiction to the applicability of bracket simplification to algebraic terms and a Revision of the application of bracket simplification to numbers.

Revision	1 "In arithmetic, brackets must 2 always be calculated first, e.g., 3 $2(5+3)/4 = 2\times8/4 = 4$
Contradiction	4 However, in algebra, the brackets 5 cannot always be simplified. 6 For instance, in $2(x+y)$, the 7 $x+y$ cannot be simplified.
Intended Message	8 So, can anything be done to 9 remove the brackets ?" *{Presentation of Distributive Law}*

Fig. 1: Distributive Law Sample Text

In this paper, we present a discourse planning mechanism which emulates a speaker's behaviour in a knowledge acquisition setting[2] in order to generate a class of rhetorical devices (RDs), denoted *Peripheral RDs*, which comprises *Contradictions*, *Revisions* and *Motivations*. Given a communicative goal to convey an *Intended Message (IM)* to a listener, these rhetorical

[1] The terms *speaker/writer* and *listener/reader* are used interchangeably in this paper.

[2] The term *knowledge acquisition setting* describes a situation where one agent transfers an intended message to another agent with the communicative goal that second agent learn the transferred information. Hence, in such a setting, intended messages are not withdrawn.

devices address beliefs presumably entertained by the listener which may affect the acquisition of the IM, but, unlike the beliefs addressed by Instantiations and Descriptions, are not directly instrumental to its comprehension [Zukerman 1990b]. Our mechanism is based on ideas introduced in [Zukerman and Cheong 1988], and it follows the tenet that Peripheral RDs are generated to invalidate anticipated impairments to a listener's comprehension process. To this effect, our mechanism departs from the existing discourse generation strategies, and adopts a predictive approach, whereby the effect of a message is simulated on a shallow model of a listener's beliefs. If an impairment to the comprehension process is anticipated by this model, the generation of a remedial rhetorical device is called for.

In the following section, we briefly consider a model of a listener's beliefs capable of predicting commonly drawn inferences. We then describe our mechanism for the generation of Peripheral RDs.

Model of a Listener's Beliefs

In order to address beliefs presumably entertained by a particular listener, we maintain an epistemological model which represents a listener's beliefs as a function of the presented material. This function accounts both for direct and indirect inferences drawn from presented messages [Zukerman 1990a].

We represent a listener's beliefs by means of a network whose nodes contain individual information items and whose links contain the relationships between the nodes (see Figure 2). The information in the network is represented at a level of detail which is consistent with the level of expertise required to learn the subject at hand, e.g., for a high-school student learning algebra, well-known concepts, such as numerical addition and subtraction, are primitive, whereas relatively new or complex concepts, such as bracket simplification, are represented in terms of more primitive concepts. The links in the network are labeled according to the manner in which they were acquired, i.e., they can either be Inferred, Told or previously Known, where Inferred links are generated by means of generally applicable *Common-sense Inference Rules*. In addition, each link is accompanied by a *Measure of Belief (MB)* between -1 and 1, akin to Certainty Factors [Buchanan and Shortliffe 1985], which represents a user's level of expertise. The nodes are labeled according to their complexity, **p** for primitive concepts and **c** for complex ones. Like links, nodes may be Inferred, Told or previously Known, and each node has a *Degree of Expertise (DE)* between 0 and 1. The DE of a c-node is a function of the DEs and MBs of its constituent nodes and links, respectively.

In this paper, we focus on technical domains, where the transmitted information typically pertains to *procedures*, *objects* and *goals*. We define a *context* as a triple composed of a procedure, an object to which it is applied, and the goal accomplished by this procedure when applied to this object (labeled c1-c4 in Figure 2), to reflect the fact that one procedure may achieve different goals when applied to different objects. For example, factoring out a common factor will only partially factorize a quadratic trinomial such as $3x^2+5x-2$, while it will completely factorize a binomial such as $3x^2+5x$.

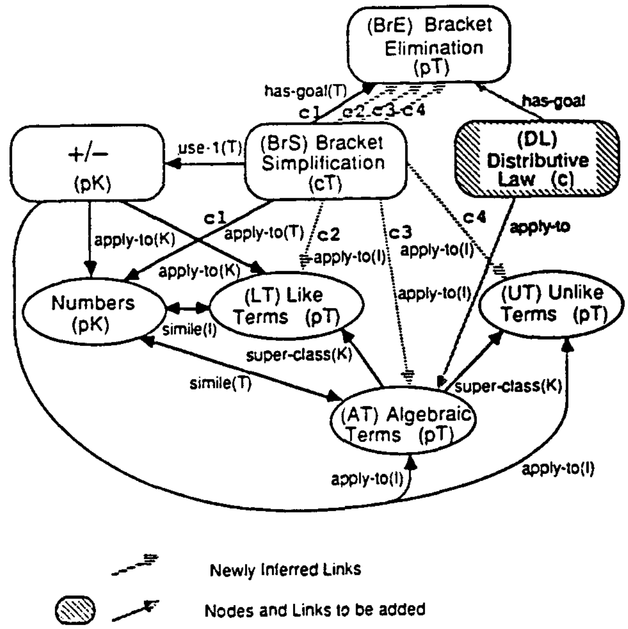

Fig. 2: Network Model of a Listener's Beliefs in High-School Algebra[3]

Our inference mechanism generates plausible inferences from links in the network by means of generally applicable Common-sense Inference Rules which portray reasoning activities such as generalization, specialization and similarity-based inference (see Figure 3). These rules are inspired by rule adaptations commonly performed by students which were studied by Matz (1982), Brown and Van Lehn (1980), Van Lehn (1983) and Sleeman (1984). In order to account for the deductive abilities of a particular type of listener, we annotate each rule with a measure of uncertainty, denoted ρ, which represents a listener's belief in the validity of a conclusion given that the evidence is certain. This measure resembles the rule strength used in ACT* [Anderson 1983]. The application of the similarity-based rule in Figure 3 to the context c1 and the link [Numbers simile AT] in Figure 2 yields the erroneous context c3.

[3] In the actual network each link may have a counterpart representing the inverse relationship. However, for clarity of presentation, only links which are relevant to our discussion are shown here.

R2 (Similarity-based Inference)
; If two objects are identified as similar, the
; applicability of a procedure to one of the objects
; can be inferred from its applicability to the other,
; accomplishing the same goal
IF { ∃ a link $[OBJ_m$ simile $OBJ_n]$ $(MB = k_{mn})$
AND
∃ a context $[PROC_a—OBJ_m—GOAL_l]$
with links $[PROC_a$ apply-to $OBJ_m]$ $(MB = k_{am})$
and $[PROC_a$ has-goal $GOAL_l]$ $(MB = k_{lam})$ }
THEN (with certainty ρ_2)
Add a context $[PROC_a—OBJ_n—GOAL_l]$
with links $[PROC_a$ apply-to $OBJ_n]$ of type I
$(MB = k_{an} = \rho_2 k_{mn} k_{am})$
and $[PROC_a$ has-goal $GOAL_l]$ of type I
$(MB = k_{lan} = \rho_2 k_{mn} k_{lam})$

Fig. 3: Similarity-based Common-sense Inference Rule

Generating Peripheral Rhetorical Devices

The generation of Peripheral RDs is performed by applying a procedure called *Impairment-Invalidate*. This procedure simulates the effect of an IM on our network model to anticipate possible impairments which may inhibit a listener's acquisition of this IM, and then proposes Peripheral RDs, such as Contradictions, Revisions and Motivations, to invalidate these impairments. To this effect, it activates three mechanisms: (1) *Propagation* of the effect of a message, (2) *Recognition* of impairments, and (3) *Selection* of Peripheral RDs. In the following subsections, we describe these mechanisms and then discuss procedure Impairment-Invalidate.

Propagation of a Message

Propagation simulates the alterations taking effect in a network representing a listener's beliefs due to inferences drawn from a message. These inferences, which are drawn by activating applicable Common-sense Inference Rules, result in changes in the MBs of existing links or addition of new links and nodes to the network. For instance, a Contradiction such as "You cannot always simplify inside the brackets of algebraic terms," presented to weaken a listener's belief in the applicability of bracket simplification to Algebraic Terms, may cause the listener to weaken his/her correct belief in the applicability of bracket simplification to Like Algebraic Terms or even to Numbers, and also to weaken his/her incorrect belief in the applicability of bracket simplification to Unlike Algebraic Terms and in the applicability of addition and subtraction to Algebraic Terms. This effect is simulated by the propagation of the Contradiction [BrS ¬apply-to AT] in the sample network in Figure 2, which produces the inferences [BrS ¬apply-to LT] and [BrS ¬apply-to UT] through the application of a specialization rule, [BrS ¬apply-to Numbers] through the application of the similarity-based rule R2 (see Figure 3), and [+/– ¬apply-to AT] through the application of a deductive inference rule. These inferences contradict the beliefs represented by existing links, thereby lowering the MBs associated with these links. However, their effect ultimately depends on the strength of the rules in question and on the MBs of the links, i.e., the impact of an inference on a weakly believed link is more pronounced than for a strongly believed one.

Recognition of Impairments

The Recognition mechanism anticipates possible impairments to a listener's comprehension process by examining changes in a network representing the listener's beliefs due to a message or an inference. It does not guarantee that a certain impairment will occur, rather, it conjectures that an impairment is likely to affect a particular link, denoted a *culprit link*.

We have characterized in terms of our network model two main types of impairments which lead to undesirable effects commonly encountered in a knowledge acquisition setting: (1) *Affect-related impairments* which elicit negative affective responses such as Confusion and Loss of Interest, and are caused by a conflict between a message or inference and a belief held by a listener; and (2) *Belief-related impairments* such as Mislearning, Insufficient Learning or Insignificant Change in a listener's knowledge status, which are caused by a discrepancy between a listener's belief in a proposition (possibly as a result of an inference) and the belief the speaker intends him/her to have in this proposition. Let us first consider the Affect-related impairments.

Confusion occurs when an inference decreases significantly a listener's confidence in a previous belief, causing a discomforting transition from a self-perception of possessing knowledge to one of increased uncertainty. In terms of our network model, Confusion takes place when the absolute value of an *Anticipated Measure of Belief (AMB)*, obtained by combining the MB of a link with the MB of an inference regarding this link, is significantly lower than the absolute value of the original MB of this link. For instance, the statement "One cannot always add algebraic terms," which yields a negative value for the MB of the link [+/– apply-to AT] in Figure 2, may trigger the erroneous inference [+/– ¬apply-to LT], which contradicts the link [+/– apply-to LT], thereby lowering its MB. The *Strength* of this impairment for a link L is defined as follows:

$$Strength(Confusion, L) = \max\{(|MB(L)| - |AMB(L)|), 0\}$$

Loss of Interest occurs when a listener who is initially motivated to acquire knowledge is presented with an IM s/he considers redundant. In terms of our model, this takes place if there exists a node B which *subsumes* a new node A, i.e., new distinguishing links incident upon A are connected to the same nodes and have MBs of compatible magnitude and sign as the corresponding links incident upon B. For this type of impairment, a culprit link is an erroneous link incident upon B representing belief, whose reversal into disbelief results in B no longer subsuming A. This situation is illustrated in Figure 2, where we try to add the node DL, representing distributive law, and the links [DL apply-to AT] and [DL has-goal BrE] to the network representing a listener's beliefs. However, the existence of the culprit link [BrS apply-to AT] supports the erroneous belief that bracket simplification is equivalent to distributive law, thereby rendering the new procedure redundant. If all the existing links participating in an impairment due to Loss of Interest are correct, no culprit link is identified.

A characterization of Belief-related impairments must take into consideration the difference between a listener's level of expertise and a level of expertise considered satisfactory. To this effect we define the *Strength* of an impairment I in link L as follows:

$$Strength(I,L) = \begin{cases} \max\{(SMB(L)-AMB(L)),0\} & \text{if } SMB(L)>0 \\ |\min\{(SMB(L)-AMB(L)),0\}| & \text{if } SMB(L)<0 \end{cases}$$

where *SMB* is a *Satisfactory Measure of Belief* representative of an adequate level of expertise with respect to a link. Its value may be obtained from a network which represents the speaker's beliefs. The relative position of Belief-related impairments and Confusion in a listener's belief space is graphically represented in Figure 4 which depicts these impairments as a function of the AMB and the previous MB of a link with $SMB>0$. The diagram for a link with $SMB<0$ is symmetric to the one in Figure 4.

Mislearning takes place when an erroneous belief with a relatively high degree of certainty is produced by an incorrect inference drawn by a listener. In terms of our network model, this takes place when the AMB of a link represents a substantial incorrect belief, and if this link existed previously, the strength of the impairment has increased, i.e., the AMB of this link is farther than its previous MB from its SMB. If the absolute value of the AMB of the link in question is higher than the absolute value of its previous MB, the listener will falsely perceive him/herself as being more proficient.

Insufficient Learning occurs when a correct inference yields a correct belief with a relatively high degree of certainty, but which still falls short of a desired degree of certainty representative of proficiency. In terms of our network model, this occurs when the AMB of a link represents a substantial correct belief, and if this link existed previously, the strength of the impairment has decreased, i.e., the AMB of this link is closer than its previous MB to its SMB.

Finally, an ***Insignificant Change*** in a listener's knowledge status occurs when an inference accomplishes a rather inconsequential change with respect to a previously non-existent link or with respect to a link

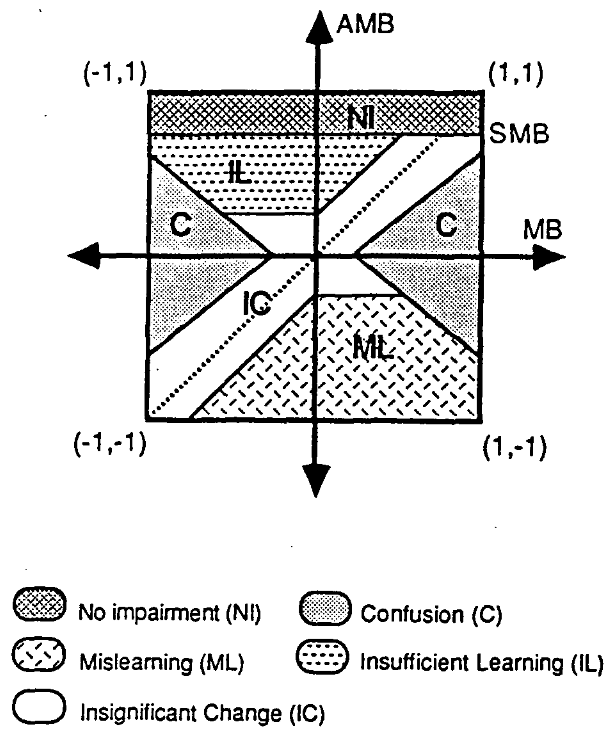

Fig. 4: Characterization of Impairments for a Link with SMB>0

with an MB representative of insufficient proficiency. This MB may represent either a correct belief or an incorrect belief.

The immediate invalidation of Affect-related impairments is essential for the smooth continuation of the knowledge acquisition process, since their persistence diverts a listener's mental resources from the task of acquiring knowledge. On the other hand, the invalidation of Belief-related impairments with respect to links which are not currently in focus may be temporarily postponed due to didactic or stylistic considerations. Therefore, although Confusion may take place in conjunction with a Belief-related impairment, the recognition and subsequent invalidation of Confusion takes precedence over the detection of this impairment (see Figure 4).

Selection of Rhetorical Devices

The Selection mechanism proposes a Peripheral RD to address a recognized impairment. The type of this rhetorical device, its wording, and its position in the final message sequence depend on the type of the impairment, the correctness and magnitude of the previous MB of the culprit link, and the SMB of this link (see Table 1). The strength of an impairment affects the need for additional explanations, such as Causal explanations and Instantiations, to convey a Peripheral RD.

Table 1: Peripheral RDs as a Function of Impairment Types and Link Values		
Impairment Type	*Link Value*	*Peripheral RD*
Loss of Interest	Correct	Motivate (add links)
	Incorrect	Contradict link
Confusion	Correct	Revise link
	Incorrect	Contradict link
Mislearning	—	Contradict inference
Insufficient Learning	—	Revise inference
Insignificant Change	High MB (Correct)	Revise link
	High MB (Incorrect)	Contradict link
	Low MB (*SMB* >0)	Revise information
	Low MB (*SMB* <0)	—

The *Link Value* column in Table 1 contains information pertaining to the previous MB and the SMB of the culprit link. This information is unnecessary when addressing Mislearning and Insufficient Learning, since these impairments are completely characterized by their type. However, for the rest of the impairments, the correctness of the previous MB is the main factor in the determination of the type of a Peripheral RD. Loss of Interest is invalidated by a Contradiction of a culprit link, if such a link is identified; otherwise, a Motivation has to generated by adding nodes and links which render the new information non-redundant, e.g., "You already know how to solve quadratic equations by completion to square. *A faster method is* ... " (details regarding the types of Motivations which are suitable for different situations and users appear in [Zukerman 1987]). The invalidation of an Insignificant Change in a link with an MB representative of lack of expertise, i.e., an MB whose absolute value is close to 0, is performed in a manner similar to the correction of complete ignorance. That is, if the SMB of this link is positive, this impairment is invalidated by a Revision of the relevant information, whereas if it is negative, i.e., representative of disbelief, no Peripheral RD is proposed, since it may be superfluous to induce disbelief with respect to a proposition which is hardly entertained by a listener.

As seen in Table 1, we distinguish between two types of Peripheral RDs according to the source of the belief being addressed, namely Peripheral RDs which address previously existing links and Peripheral RDs which address current inferences. This distinction is not always reflected in the English realization of a rhetorical device, rather, it may affect the Meta Comments which accompany it and its position in the final message sequence. For instance, the propagation of the Contradiction [BrS ¬apply-to AT] may yield the erroneous inference [BrS ¬apply-to LT] which in turn may cause an impairment in the link [BrS apply-to LT]. If the anticipated impairment is Confusion, it will be invalidated by a Revision of this link, whereas if it is Mislearning, it will be invalidated by a Contradiction of the erroneous inference. The most succinct realization of both Peripheral RDs is the sentence "You can always simplify bracketed Like Terms." However, the Revision of the previous belief may appear either before or after the above Contradiction and include a Meta Comment which states the source of this belief, e.g., "*As we saw in Section 7,* you can always simplify bracketed Like Terms"; while the Contradiction of the erroneous inference would usually appear after the above Contradiction and would be accompanied by a Meta Comment which indicates a violation of an expectation established by the inference, e.g., "Bracket simplification does not always apply to algebraic terms, *but* it always applies to Like Terms." The Peripheral RDs proposed by our Selection mechanism contain sufficient information to support the generation of these types of Meta Comments by means of mechanisms such as the ones presented in [Zukerman and Pearl 1986] and [Zukerman 1989].

Procedure Impairment-Invalidate

Impairment-Invalidate is activated with one argument which contains an IM. It returns a *Message-List* composed of the IM and the Peripheral RDs which were proposed to invalidate the impairments anticipated as a result of this message.

```
Procedure Impairment-Invalidate(Message)
Message-List ← Message
Peripheral-RDs ,Inferences ← nil
FirstImp ← (Message ,T_M ,S_M ,L_M),
      where (T_M ,S_M ,L_M) ← Recognize (Message)
If FirstImp Then
   Message-List ←
         Append (Message-List ,Select (FirstImp))
For each m ∈ Message-List do
   Inferences ← Merge (Inferences ,Propagate (m))
endfor
Impairments ← {(i ,T_i ,S_i ,L_i) | (i ∈ Inferences) ∧
                    ((T_i ,S_i ,L_i) ← Recognize (i))}
While Impairments do
   MaxImp ← (I ,T_I ,S_I ,L_I),
      where {(I ,T_I ,S_I ,L_I) ∈ Impairments ∧
            Ranking (I ,T_I ,S_I ,L_I) =
               max_{(j ,T_j ,S_j ,L_j) ∈ Impairments} {Ranking (j ,T_j ,S_j ,L_j)}}
   Impairments ← Impairments − MaxImp
   RD_MaxImp ← Select (MaxImp)
   Peripheral-RDs ←
         Append (Peripheral-RDs ,RD_MaxImp)
   Secondary-Inferences ← Propagate (RD_MaxImp)
   Secondary-Effects ← Secondary-Inferences ∩
            {j | (j ,T_j ,S_j ,L_j) ∈ Impairments}
   Impairments ← {Impairments −
         {(i ,T_i ,S_i ,L_i) | i ∈ Secondary-Effects}} ∪
         {(i ,T_i ,S_i ,L_i) | (i ∈ Secondary-Inferences) ∧
                  ((T_i ,S_i ,L_i) ← Recognize (i))}
endwhile
Message-List ←
      Append (Message-List ,Peripheral-RDs)
```

Table 2: Inferences and Possible Impairments after Propagation of the IM and the Contradiction in the Sample Network			
Message	*Rule*	*Inference*	*Possible Impairment*
IM [DL apply-to AT has-goal BrE]	Similarity	[DL apply-to Numbers has-goal BrE]	*Insufficient Learning*
	Specialization	[DL apply-to LT has-goal BrE]	
	Specialization	[DL apply-to UT has-goal BrE]	
Contradiction [BrS ¬apply-to AT]	Similarity	[BrS ¬apply-to Numbers]	*Confusion/Mislearning/ Insignificant Change*
	Specialization	[BrS ¬apply-to LT]	
	Specialization	[BrS ¬apply-to UT]	*Confusion/Insuf. Learning/ Insignificant Change*
	Deduction	[+/− ¬apply-to AT]	

We distinguish between two main stages of this procedure: (1) The preliminary stage (lines 1-4) which determines the need for a Peripheral RD related to the input message, and (2) The iterative stage (lines 5-17) which ascertains the need for Peripheral RDs pertaining to subsequent inferences.

In the preliminary stage, procedure *Recognize* is applied in order to determine whether the IM is likely to cause an impairment. If this is the case, Recognition returns a triple (T,S,L), where T contains the *Type* of the impairment, S its *Strength*, and L indicates whether the belief represented by the link in question is correct or incorrect. Otherwise, it returns *nil* and no impairment is predicted. If an impairment was anticipated, procedure *Select* is activated to propose a Peripheral RD. In our example, the Selection mechanism invalidates the impairment responsible for the Loss of Interest by means of the Contradiction [BrS ¬apply-to AT] which induces disbelief in the link [BrS apply-to AT]. Upon completion of the preliminary stage, the proposed Peripheral RD together with the IM form the Message-List, which constitutes the input to the next phase of the impairment invalidation procedure.

Both the IM and the proposed Peripheral RD cause modifications in a listener's degree of belief in the addressed links. These modifications in turn may lead to changes in his/her beliefs in other links. Hence, in order to prevent impairments due to inferences drawn from these messages, additional Peripheral RDs may be called for. Furthermore, Peripheral RDs may be required in order to invalidate an Insignificant Change in links related to the IM which have MBs representative of insufficient proficiency, i.e., to attain a satisfactory degree of expertise with respect to these links. To determine the need for additional Peripheral RDs, *Propagate* produces inferences from each message in Message-List. During this process, inferences from different messages which affect the same link are merged into one inference with a combined effect (lines 5-7). Recognition then ascertains the attributes of the impairments which are likely to be caused by these inferences (line 8). For instance, the Propagation of the input message [DL apply-to AT has-goal BrE] and the Contradiction [BrS ¬apply-to AT] in the sample network in Figure 2 may yield the inferences in Table 2. In principle, each of these inferences may be responsible for an impairment. However, as stated above, the effect of these inferences ultimately depends on the ρ-s of the rules applied in the propagation process and on the MBs of the affected links. In our present discussion, we assume that Confusion was recognized in the links [BrS apply-to LT] and [BrS apply-to Numbers].

Based on Gricean maxims of cooperative conversation (Grice 1975), we propose to generate a minimal number of rhetorical devices to invalidate impairments occurring concurrently in a number of links. To this end, we iterate over the set of impairments, selecting at each stage the culprit link with the highest ranking impairment (line 10). The impairments are ranked according to their type and strength, where impairments causing Confusion are ranked higher than Belief-related impairments. A Peripheral RD is then proposed to invalidate the impairment in the selected link, and its effect is propagated (lines 12-14). Once a Peripheral RD has been generated to address a particular link, further inferences drawn during the same activation of Impairment-Invalidate no longer affect this link. For each iteration, the set of impairments is updated by merging the inferences responsible for the previously recognized impairments with the inferences resulting from the most recent propagation, and reapplying the Recognition process (lines 15-16). In this manner, a low-ranking impairment in a given link may be spontaneously invalidated by an inference resulting from a Peripheral RD generated to invalidate an impairment in another link. For example, if the impairment in the link [BrS apply-to Numbers] ranks higher than the impairment in the link [BrS apply-to LT], a Revision is generated to invalidate the impairment in the former link. The propagation of this Revision may invalidate the Confusion with respect to the latter link. If this Revision does not cause further impairments, the procedure terminates returning the following messages:

IM	[DL apply-to AT has-goal BrE]
Contradiction	[BrS ¬apply-to AT]
Revision	[BrS apply-to Numbers has-goal BrE]

The generation of belief Revisions tends to inhibit further impairments, since they reinforce beliefs which are likely to be consistent with the rest of a listener's beliefs, while the generation of belief Contradictions tends to foster impairments, since they disagree with existing beliefs. However, in a knowledge acquisition setting, misconceptions are generally not allowed to pile up, hence Contradictions to existing beliefs are expected only during the initial stages of the propagation process. Therefore, although in principle our algorithm may iterate indefinitely, in practice, impairments should no longer be detected after a few iterations, and the process should halt after proposing a few Peripheral RDs. This expectation is confirmed by tests run with our sample network (see next section). However, this network is rather small, and there may be situations in which the number of Peripheral RDs proposed by our mechanism will have to be restricted due to stylistic or pedagogical considerations. In such cases, our mechanism will have to be adjusted to invalidate only a subset of the recognized impairments such as the highest ranking impairments or impairments in links which are closest to the IM.

Evaluation

Procedure Impairment-Invalidate has been implemented in a system called WISHFUL, which was run on several instances of the network in Figure 2. These instances featured MBs which represented three types of students: (1) Competent — with high MBs associated with correct beliefs, (2) Average — containing some incorrect beliefs and medium-range MBs associated with correct beliefs, and (3) Mediocre — with very low MBs associated with most beliefs. The initial values assigned to the MBs of the links and the ρ-s of the Common-sense Inference Rules yielded Peripheral RDs which were compatible with rhetorical devices people would generate under similar circumstances, and the response time was instantaneous for these rather small networks. As expected, changes in the ρ-s of the Common-sense Inference Rules resulted in variations in the generated RDs, with additional Contradictions due to Mislearning being generated as the ρ-s of unsound rules increased, and additional Revisions due to Insufficient Learning as the ρ-s of sound rules decreased. In addition, the number and strength of the recognized impairments increased as the ability of the students being represented in a network decreased, indicating that additional explanations would be required to convey a message to the poorer students. In particular, for the networks representing competent students, either Revisions of inferences due to Insufficient Learning or no Peripheral RDs were proposed (depending on the ρ-s of the Common-sense Inference Rules); for the networks representing average students, Contradictions and Revisions were proposed in a manner similar to the explanations in this paper; and for the networks representing mediocre students most of the proposed Peripheral RDs were Revisions of information.

Our mechanism has also proven successful as an analytical tool. It accounts for the presence of Peripheral RDs in over twenty texts in a variety of domains, ranging from expert domains (Telecommunications, Cognitive Science and Linguistics) through intermediate ones (high-school Algebra, Data Structures, Lisp and introductory Chess) to novice domains (Childcraft Encyclopedia and Dr. Spock's Baby and Child Care).

Limitations and Future Research

Our mechanism proposes rhetorical devices under the assumption that after it has done "its best," the listener's beliefs addressed by the discourse will be modified in the desired direction, although not necessarily to the desired extent. This is a valid assumption for discourse generation, since one can not say more than one knows. However, for this mechanism to generate effective discourse on a continued basis, the model of the listener's beliefs must be updated by an independent assessment of the his/her understanding.

In addition, our mechanism must be implemented on different domains and larger networks to test both its response time and the adequacy and number of the proposed rhetorical devices. As stated above, this may reveal the need to adjust procedure Impairment-Invalidate to enable it to control the number of the proposed Peripheral RDs.

At present, research is in progress to extend the impairment invalidation paradigm to the generation of other types of rhetorical devices, such as Descriptions, Instantiations and Causal explanations, and to devise an algorithm to sort the generated messages according to rhetorical considerations (Zukerman 1990b). In addition, an alternative representation for MBs which keeps track of the sources of an inference is being considered. Further research is required to recognize and rectify possible misconceptions in the Common-sense Inference Rules, and to characterize conditions for the generation of rhetorical devices which satisfy a number of communicative goals. Finally, the effect of the knowledge representation and the rules of inference on the types of the proposed rhetorical devices needs to be further investigated.

Conclusion

This paper offers a text planning mechanism which supports the generation of explanations tailored to particular types of users. Our mechanism generates Peripheral RDs which help convey an intended message by anticipating and preventing potential impairments to a listener's comprehension process. To this effect, it characterizes these impairments in terms of a model of a listener's beliefs and inferences, and simulates a listener's comprehension process on this model. Clearly, this process does not dispense with the need to interact with a listener, but it addresses commonly occurring impairments, thereby focusing the interaction. Furthermore, it is envisioned that the impairment invalidation mechanism will become a useful tool to guide the generation of cogent responses to user follow-up queries, as it can point to issues which are potentially troublesome.

References

[Anderson, 1983] J.R. Anderson. *The Architecture of Cognition*, Harvard University Press, Cambridge, Massachusetts.

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

[Brown and Van Lehn, 1980] J.S. Brown and K. Van Lehn. Repair Theory: A Generative Theory of Bugs in Procedural Skills. In *Cognitive Science* 4, pp. 379-426.

[Buchanan and Shortliffe, 1985] B.G. Buchanan and E.H. Shortliffe. *Rule-Based Expert Systems — The MYCIN Experiments of the Stanford Heuristic Programming Project*, Addison-Wesley Publishing Company.

[Grice, 1975] H.P. Grice. Logic and Conversation. In P.J. Cole and J.L. Morgan (Eds.), *Syntax and Semantics, Volume 3: Speech Acts*, Academic Press, pp. 41-58.

[Hovy, 1987] E.H. Hovy. *Generating Natural Language under Pragmatic Constraints*. Doctoral Dissertation, Computer Science Department, Yale University, New Haven, Connecticut.

[Hovy, 1988] E.H. Hovy. Planning Coherent Multisentential Text. In *Proceedings of the Twenty-Sixth Annual Meeting of the Association for Computational Linguistics*, State University of New York, Buffalo, New York.

[Kukich, 1983] K. Kukich. *Knowledge-Based Report Generation: A Knowledge-Engineering Approach to Natural Language Report Generation*. Doctoral Dissertation, The Interdisciplinary Department of Information Science, University of Pittsburgh, Pennsylvania.

[Lynch *et al.*, 1979] B.J. Lynch, R.E. Parr and H.M. Keating. *Maths 8*, Sorrett Publishing.

[Matz, 1982] M. Matz. Towards a Process Model for High School Algebra Errors. In D. Sleeman and J.S. Brown (Eds.), *Intelligent Tutoring Systems*, London: Academic Press, 1982, pp. 25-50.

[McKeown, 1985] K.R. McKeown. Discourse Strategies for Generating Natural Language Text. In *Artificial Intelligence* 27, pp. 1-41.

[Moore and Swartout, 1989] J.D. Moore and W.R. Swartout. A Reactive Approach to Explanation. In *IJCAI-11 Proceedings, International Joint Conference on Artificial Intelligence*, pp. 1504-1510.

[Paris, 1988] C.L. Paris. Tailoring Object Descriptions to a User's Level of Expertise. In *Computational Linguistics*, Volume 14, Number 3, September 1988, pp. 64-78.

[Sleeman, 1984] D. Sleeman. Mis-Generalization: An Explanation of Observed Mal-rules. In *Proceedings of the Sixth Annual Conference of the Cognitive Science Society*, pp. 51-56.

[Van Lehn, 1983] K. Van Lehn. Human Procedural Skill Acquisition: Theory, Model and Psychological Validation. In *AAAI-83 Proceedings, American Association for Artificial Intelligence*, August 1983, pp. 420-423.

[Zukerman and Pearl, 1986] I. Zukerman and J. Pearl. Comprehension-Driven Generation of Meta-technical Utterances in Math Tutoring. In *AAAI-86 Proceedings, American Association for Artificial Intelligence*, pp. 606-611.

[Zukerman, 1987] I. Zukerman, Goal-based Generation of Motivational Expressions in a Learning Environment. In *AAAI-87 Proceedings, American Association for Artificial Intelligence*, July 1987, pp. 327-331.

[Zukerman and Cheong, 1988] I. Zukerman and Y.H. Cheong. Impairment Invalidation: A Computational Model for the Generation of Rhetorical Devices. In *Proceedings of the International Computer Science Conference '88: Artificial Intelligence, Theory and Applications*, pp. 294-300.

[Zukerman, 1989] I. Zukerman. Expectation Verification: A Mechanism for the Generation of Meta Comments. In *Proceedings of the 11th Conference of the Cognitive Science Society*, Ann Arbor, Michigan, pp. 498-505.

[Zukerman, 1990a] I. Zukerman. A Predictive Approach to the Generation of Rhetorical Devices. To appear in *Computational Intelligence*, Vol 6, No. 1.

[Zukerman, 1990b] I. Zukerman. Anticipating a Listener's Response in Text Planning. To appear in *Advances in Artificial Intelligence, Natural Language and Knowledge-based Systems*, Springer Verlag.