Generating Concise Discourse that Addresses a User's Inferences*

Ingrid Zukerman and Richard McConachy
Department of Computer Science
Monash University
Clayton, Victoria 3168
AUSTRALIA

Abstract

In this paper, we describe a content planning mechanism which takes into consideration a user's possible inferences in order to generate the most concise discourse that achieves a given communicative goal. The consideration of a user's inferences results in the addition of information that addresses erroneous inferences, and the omission of easily inferred information. Given a communicative goal, our mechanism applies inference rules in backward reasoning mode to plan rhetorical devices that achieve this goal, and subsequently in forward reasoning mode to conjecture the effect of these rhetorical devices on the user's beliefs. This process results in one or more sets of rhetorical devices which achieve the intended communicative goal. Each set is then minimized, and the most concise among these sets is selected. These ideas have been implemented in a system called WISHFUL, which generates explanations about concepts.

1 Introduction

Traditional Natural Language Generation (NLG) systems, e.g., [Hovy, 1988; Moore and Swartout, 1989; Cawsey, 1990], operate under the assumption that a hearer will only make direct inferences from a piece of discourse. This assumption causes two problems:

(1) *Possible erroneous indirect inferences from the discourse are not addressed.* For example, a possible indirect inference from the statement "plants photosynthesize" is that fungi also photosynthesize. Even though this inference is wrong, it would not be addressed by current systems.

(2) *The resulting text is overly explicit* due to the fact that all the information to be conveyed is stated. For instance, given the goal of conveying the propositions [Mary went-to cinema] and [Mary saw film], existing NLG systems would generate a sentence for each proposition, yielding discourse such as "Mary went to the cinema and she saw a film." However, this discourse is overly explicit, since, unless told otherwise, most people infer from the first sentence that Mary saw a film.

These problems have been considered separately in previous research. The first of these problems has been addressed in [Joshi *et al.*, 1984; Zukerman, 1990] in the context of adding explanations to expert responses, and complementing planned discourse with Contradictions and Revisions, respectively. [Horacek, 1991; Lascarides and Oberlander, 1992] address the second problem by omitting from planned explanations information that is easily inferred from the discourse.

In this paper, we present an integrated content planning mechanism which addresses both of the above problems. Our mechanism generates *Rhetorical Devices (RDs)*, where each RD is composed of a rhetorical action, such as Assert, Negate or Instantiate, applied to a proposition. To this effect, our mechanism models a user's inferences by means of inference rules which relate RDs to beliefs (Section 2). These rules are applied in two different ways during the discourse planning process: *forward reasoning* and *backward reasoning*.

Forward reasoning reasons from RDs to their possible effects. For instance, the application of a generalization inference rule to the Assertion "kangaroos are indigenous of Australia" conjectures that the hearer will conclude that marsupials are indigenous of Australia.

Backward reasoning reasons from a communicative goal to the RDs that may be used to accomplish it. For instance, the concept of a stack may be conveyed to a student by means of a Definition, an Analogy (say to a stack of plates in a cafeteria), an Instantiation, or a combination of these RDs. This reasoning mechanism has been widely used in NLG systems, e.g., [Hovy, 1988; Moore and Swartout, 1989; Cawsey, 1990]. In particular, [Moore and Swartout, 1989; Cawsey, 1990] have encoded Analogies and Instantiations into discourse planning operators. However, this encoding does not represent explicitly the inferences that allow a hearer to deduce a belief from a rhetorical device. In addition, according to these operators, the user acquires only the intended beliefs, and no unintended ones.

Our content planning mechanism has been implemented in a system called WISHFUL which generates concise discourse while at the same time addressing a user's possible erroneous inferences. Our mechanism first applies backward reasoning in order to determine which RDs are suitable for communicating a particular piece of information, and then applies forward reasoning in order to conjecture the effect of these RDs on a user's beliefs. If a wrong inference is conjectured from an RD, then an additional RD may be required to address this inference. If a correct inference is conjectured, then a previously planned RD which conveys the inferred information may no longer be required. Owing to these inter-

*This research was supported in part by grant A49030462 from the Australian Research Council.

actions between the inferences from the RDs generated to convey different information items, an optimization process is applied in order to achieve a global minimum of the number of generated RDs (Section 3.4).

In the next section, we discuss our model of the student's beliefs and inferences. In Section 3, we describe the operation of the content planner. We then discuss our results and present concluding remarks.

2 The Student Model

Our student model represents (1) a student's profile, (2) his/her beliefs, and (3) his/her inferences.

Profile: [Sleeman, 1984] observed that good students retain more correct conclusions than incorrect ones, while the opposite happens for mediocre students. In addition, we have observed that good students are more certain of their conclusions than mediocre students. In order to model these behaviours, we maintain profiles of different types of students. The profile attributed to a student determines the correctness and strength of the initial beliefs in the model of the student, and the degree of belief in the conclusions drawn by sound and unsound inferences. For example, the profile of a mediocre student is characterized by weak convictions with respect to both facts and inference rules, and lack of discrimination between correct and incorrect beliefs, and between sound and unsound inferences. In the current implementation of WISHFUL we maintain five profiles which range from EXCELLENT to BAD.

Beliefs: Since it is easier to make broad assessments rather than pinpoint numerical assessments with respect to a student's beliefs, we represent a student's conjectured beliefs by means of the following *qualitative belief states* [Bonarini *et al.*, 1990]: {*BELIEVED, RATHER BELIEVED, CONTRADICTORY, UNKNOWN, RATHER DISBELIEVED, DISBELIEVED*}.

Inferences: Our mechanism uses inference rules to model three types of inferences made by the user: (1) *Direct inferences*, (2) *Indirect inferences*, and (3) *Uniqueness implicatures*.

Direct inferences reproduce directly the content of the discourse. The *abstract-understand* inference rule assesses the likelihood that a hearer will understand a statement by means of a direct inference. This likelihood is influenced by the complexity and abstractness of the information in the statement and by the hearer's ability to understand abstract explanations, such as stand-alone descriptions or definitions. This ability in turn is represented in the hearer's profile.

Indirect inferences produce inferences that add information to what was said. These inferences are not always sound. The indirect inference rules considered in our model are based on the ones described in [Zukerman, 1990], e.g., *generalization*, *specialization* and *similarity*. The likelihood that a hearer will acquire a belief through an indirect inference depends on his/her confidence in the corresponding inference rule and on the strength of the beliefs which participate in the inference process.

Finally, given a proposition $P(O)$, a **uniqueness implicature** licenses the inference that P is true *only* with respect to O. For example, upon hearing the statement "Joe has one leg," most people will infer that Joe has one leg *only* [Hirschberg, 1985].

3 Operation of the Content Planner

Our content planner receives as input a *concept* to be conveyed to the hearer, e.g., Distributive Law, a list of *aspects* that must be conveyed about this concept, e.g., steps and domain, and an *attitude*, which may be either *cautious* or *daring*. A cautious attitude demands that the information relevant to the given aspects be *BELIEVED* by the user, while a daring attitude is satisfied with *RATHER BELIEVED*. The output of the content planner is a set of RDs which conveys the intended aspects of the given concept. Typically, a set of RDs generated with the cautious attitude contains more RDs and more complex RDs than a set of RDs produced with the daring attitude. That is, the explanations generated with the cautious attitude are typically longer and more thorough than those generated with the daring attitude.

Our content planner performs the following steps: First it determines the information to be presented based on the given aspects **(Step 1)**. Next it takes into consideration a user's possible inferences in order to propose alternative sets of RDs which convey this information **(Step 2)**. However, it is possible that the user does not understand the concepts mentioned in a particular RD well enough to understand this RD. Therefore, the generation process is applied recursively with a revised attitude and new aspects with respect to the concepts mentioned in the proposed sets of RDs **(Step 3)**. If necessary, this process generates alternative sets of RDs which convey these concepts. Thus, Steps 1-3 of WISHFUL yield several alternative sets of RDs, where each set contains enough information to convey the intended concept. However, owing to the interactions between the inferences from the RDs in each set, it is possible that some of the proposed RDs are no longer necessary. In order to remove the redundant RDs from each set of RDs, WISHFUL applies an optimization process to each set **(Step 4)**. It then selects the set with the least number of RDs from the resulting sets **(Step 5)**.

Throughout this section, we use the following sample input to illustrate the operation of the content planner: **(**Bracket-Simplification**,** {*steps,domain*}**, daring)**. In this input, the communicative goal is for the hearer to *RATHER BELIEVE* the information relevant to the domain and steps of Bracket Simplification.

3.1 Deciding which Information to Present

In this step, WISHFUL produces a list of propositions that must be conveyed in order to satisfy a given communicative goal with respect to the specified aspects of a given concept. To this effect, WISHFUL first retrieves from a knowledge base the propositions relevant to the given aspects. For instance, in order to satisfy the aspects in our sample input, the propositions in Table 1 must be known by the hearer[1].

Aspect	Propositions
steps	p_1: [Bracket-Simplification use-1 +/−] p_2: [Bracket-Simplification use-2 ×]
domain	p_3: [Bracket-Simplification apply-to Like-Terms] p_4: [Bracket-Simplification apply-to Numbers]

Table 1: *Propositions Relevant to steps and domain*

[1]The relationships use-1 and use-2 indicate the temporal ordering of a mathematical operation.

The retrieved list of propositions is then refined based on consultation with our model of the hearer's beliefs. The propositions already known by the hearer are filtered out, and propositions which address the hearer's misconceptions are added. Thus, in our example, if the hearer is presumed to believe correctly that [Bracket-Simplification use-2 ×] and [Bracket-Simplification apply-to Numbers], and incorrectly that [Bracket-Simplification apply-to Algebraic-Terms], the final list of propositions will be as shown in Table 2.

p_1: [Bracket-Simplification use-1 +/−]
p_3: [Bracket-Simplification apply-to Like-Terms]
p_5: [Bracket-Simplification ¬(always) apply-to Algebraic-Terms]

Table 2: *Propositions to be Conveyed*

3.2 Proposing Rhetorical Devices

In this step, the content planner activates the procedure *Propose-RDs* to propose alternative sets of RDs that convey the propositions produced in the previous step. Our procedure takes into consideration the inferences a hearer is likely to make from the RDs in each alternative in order to conjecture the effect of this alternative on the hearer's beliefs. The operation of our procedure is based on the tenet that people draw immediate inferences while processing a piece of discourse, but make farther reaching inferences only after the entire discourse has been processed. In order to address these immediate inferences, our procedure applies one round of inference rules to each RD in an alternative. For instance, given the RD **Assert** [Bracket-Simplification apply-to Like-Terms], the generalization rule infers incorrectly that [Bracket-Simplification apply-to Algebraic-Terms], and the similarity rule yields the correct inference that [Bracket-Simplification apply-to Numbers].

Procedure **Propose-RDs**($\{propositions\}$)
1. If $\{propositions\}=\emptyset$ Then return(*nil*).
2. Get a proposition$\in\{propositions\}$, and remove it from $\{propositions\}$.
3. **Backward reasoning**: propose a set {RD}, where each RD in the set conveys this proposition.
4. **Forward reasoning**: $\forall RD_i \in \{RD\}$ Do (each RD_i is the root of a different alternative)
   4.1 Draw inferences from RD_i.
   4.2 Append to $\{propositions\}$ any new proposition inferred from RD_i.
   4.3 Create a link between RD_i and each of the propositions $p_{j_1}, \ldots, p_{j_n}$ that are affected by it.
   4.4 Assign to each link a weight w_{i,j_k} equal to the effect of RD_i on the belief in p_{j_k}.
   4.5 RDs ← cons(RD_i, **Propose-RDs**($\{propositions\}$)).

Propose-RDs is a recursive procedure, where each cycle of the recursion generates alternative RDs that convey one proposition. The alternatives generated during a cycle constitute the basis for the next cycle of the recursion. In each cycle, Propose-RDs first performs backward reasoning and then forward reasoning. During backward reasoning (Step 3), the system applies inference rules to propose alternative RDs that convey a proposition. In this process, the system also takes into consideration the attributes of the information to be conveyed and the information in the system's knowledge base and in the model of the user's beliefs. During forward reasoning (Step 4), the inference rules are applied to each proposed RD and to our model of the user's beliefs in order to postulate the user's possible inferences from this RD. Any new propositions affected by these inferences are then added to the list of propositions to be conveyed (Step 4.2), and are addressed in subsequent calls to Propose-RDs if necessary (Step 4.5). In principle, if the system keeps postulating that the user will infer new erroneous beliefs, the algorithm will not terminate. However, in conversations in general, and in tutoring situations in particular, the transfer of information typically takes place at the horizon of a body knowledge that is mutually believed by the conversational partners. Thus, in practice, the inferences from the presented information reach this body of knowledge quickly, thereby leading to a bounded inference process.

The RDs relevant to our domain are (A) *Assertion*; (N) *Negation*; (M) *Mention*, which acknowledges a correct belief held by the hearer, e.g., "*In addition to Numbers,* Bracket Simplification applies to Like Terms"; (I) *Instantiation*; and (I^+) *Expanded Instantiation*, which is an Instantiation annotated with brief comments, e.g.,

$$\underbrace{2(3x+5x)}_{\text{add terms in brackets}} = \underbrace{2(8x)}_{\text{multiply}} = 16x.$$

Assertions and Negations may be complemented either with Instantiations or with Expanded Instantiations in order to have a stronger effect on a hearer's beliefs. However, the suitability of a combination of RDs depends on the aspect being conveyed and on the RD being complemented. For example, when conveying the steps of a procedure, A, $A+I$ and $A+I^+$ are possible candidates. However, when conveying the domain of a procedure, only A and $A+I$ are applicable, since what is being conveyed is that the procedure applies to some domain, not how this is done. Finally, when refuting the domain of a procedure, only N and $N+I^+$ are applicable, since when illustrating that a procedure does not apply to an object, it is essential to explain where the procedure fails.

Figure 1 depicts a partial trace of the alternatives generated by Propose-RDs in order to convey the propositions $\{p_1, p_3, p_5\}$ in Table 2. The backward reasoning process appears in boldface, and the forward reasoning process in roman font and italics. After each backward reasoning step, the state of the resulting alternative is presented. A state contains (1) the RDs proposed so far; (2) the propositions affected by these RDs; (3) the degree of belief in these propositions as a result from the RDs generated so far (*BELIEVED (B), RATHER BELIEVED (RB), CONTRADICTORY (C), UNKNOWN (U), RATHER DISBELIEVED (RD)* and *DISBELIEVED (D)*); and (4) the effect of the inferences from the RDs on the user's beliefs in these propositions. This effect is either a quantitative increment or reduction in belief. However, for clarity of presentation, we represent only the trend of this effect, i.e., + (increment), − (reduction) and ± (two or more inferences with opposite effects).

In the first cycle in Figure 1, the backward reasoning stage proposes the following RDs to convey proposition p_1 ([Bracket-Simplification use-1 +/−]): (1) A, (2) $A+I$, and (3) $A+I^+$. Since Bracket Simplification applies to Numbers and to Like Terms, an Instantiation may be generated for each of these domains. Thus, alternatives (2) and (3) spawn two options each: one where the Instantiation is performed with respect to Numbers, and another where the Instantiation is performed with respect to Like Terms. However, for the sake of brevity, we limit our discussion to the latter. In the forward reasoning stage, p_1 is directly inferred from all three al-

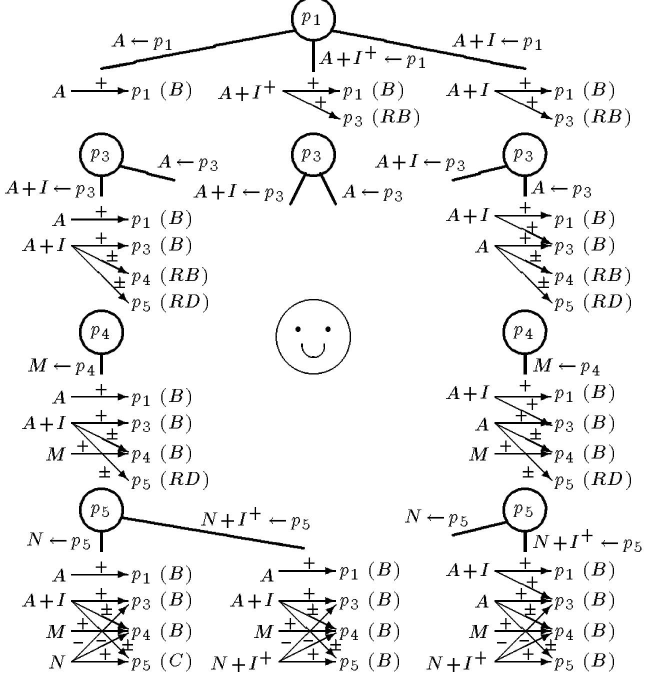

Figure 1: *Partial Trace of the Alternatives Generated by Propose-RDs to Convey* $\{p_1, p_3, p_5\}$

ternatives. In addition, a generalization from $A{+}I$ and $A{+}I^+$, which are instantiated with respect to Like Terms, conveys p_3 ([Bracket-Simplification apply-to Like-Terms]).

In the second cycle, the backward reasoning stage proposes the RDs A and $A{+}I$ to convey proposition p_3 (recall that $A{+}I^+$ is not an option in this case). In the forward reasoning stage, the following inferences ensue from both options: (1) a direct inference which yields p_3, (2) a generalization which supports the user's incorrect belief that [Bracket-Simplification apply-to Algebraic-Terms] ($\neg p_5$), (3) a similarity-based inference which supports the user's correct belief that [Bracket-Simplification apply-to Numbers] (p_4), and (4) a uniqueness implicature which infers that Bracket Simplification applies *only* to Like Terms, thereby contradicting inferences (2) and (3). Note that proposition p_4, which was previously removed from the list of propositions to be conveyed (Section 3.1), is added to this list, since it is affected by a forward inference. In addition, although p_3 is now *BELIEVED* in both alternatives, in the alternative on the right hand side, this belief is a relatively small increment from *RATHER BELIEVED*, which was obtained in the first cycle.

In the third cycle, the backward reasoning stage proposes the RD M to convey p_4 ([Bracket-Simplification apply-to Numbers]). As stated above, a Mention acknowledges a correct belief held by the hearer, rather than presenting new information. Since inferences from this belief were drawn at the time the belief was acquired, indirect inferences from the Mention should not be considered. Hence, in the forward inference stage, only p_4 is inferred from this RD.

Finally, in the fourth cycle, the backward reasoning stage proposes the RDs N and $N{+}I^+$ to convey proposition p_5 ([Bracket-Simplification $\neg$(always) apply-to Algebraic-Terms]) (as stated above, $N{+}I$ is not an option in this case). In the forward reasoning stage, the following inferences ensue from both options: (1) a direct inference which yields p_5, (2) a specialization which contradicts p_3, and (3) a similarity-based inference which contradicts p_4. Note that $N{+}I^+$ reverses the belief in p_5 from *RATHER DISBELIEVED* to *BELIEVED*, while N takes it only to *CONTRADICTORY*.

An alternative can be discarded based on insufficient belief in one or more of the propositions to be conveyed only after RDs for all the propositions have been generated. This is because an inference from an RD proposed late in the process may support a marginal belief resulting from an earlier RD. For instance, upon completion of Propose-RDs, the belief in p_5 has not reached an acceptable level in the leftmost alternative in Figure 1. Hence, this alternative is discarded.

As seen in this process, RDs may be generated at a later stage to convey propositions that have already been conveyed by indirect inferences from earlier RDs. For instance, in cycle 2, right hand column, RDs are generated to convey p_3 even though it has been inferred from the $A{+}I$ which was generated to convey p_1. The rationale for this policy is twofold. Firstly, it is possible that an inference from a later RD will lower the belief in p_3, thus requiring that we generate RDs that convey p_3 after all. But more importantly, an RD that conveys p_3 may be necessary to achieve optimality. For example, this happens when an RD that conveys p_3 directly is also capable of conveying every other proposition in our list. In this case, if the RDs that convey p_3 are not generated, the optimal solution will be missed.

The alternatives produced by Propose-RDs may contain redundant RDs. However, at this stage of the content planning process it is premature to minimize the number of RDs in each alternative, since alternatives that appear promising at this stage may be expanded by RDs that explain the concepts mentioned in these alternatives (Step 3 of WISHFUL). Therefore, the minimization process must be performed only after all the RDs have been generated for each alternative.

3.3 Conveying the Concepts in an Alternative

This step receives as input the sets of RDs produced in the previous step. For each of these alternatives, WISHFUL ascertains that the hearer understands the concepts mentioned in its RDs well enough to understand these RDs. To this effect, it performs the following actions for each concept in an alternative: (1) it determines the aspects of the concept which are relevant to the understanding of the RDs which contain this concept, (2) it determines an attitude for conveying this concept, and (3) it re-activates Steps 1-3 of WISHFUL to generate RDs that convey this concept.

The aspects a hearer must know about a concept in order to understand an RD which contains this concept depend on (1) the main predicate of the propositional part of the RD, and (2) the role of the concept with respect to this predicate. For example, in order to understand the RD **Assert** [Bracket-Simplification apply-to Like-Terms], the hearer must know what Like-Terms are and what they look like. Hence, the system returns the aspects *membership-class* and *structure*. Since a concept may appear in more than one RD, this process is repeated with respect to all the RDs which mention a particular concept.

The determination of an attitude for conveying the newly determined aspects of a concept is based on the relevance of this concept to the original concept given to the system. That is, the more relevant the concept is to the original concept, the better it should be known

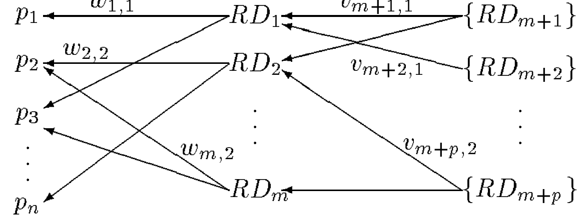

Figure 2: *A Set of RDs which Conveys* $\{p_1, \ldots, p_n\}$

by the hearer, and vice versa. This consideration is implemented by changing an initially cautious attitude to daring as the recursive calls to Steps 1-3 of WISHFUL become deeper. A daring attitude is not changed, since it already requires a low level of expertise.

3.4 Minimizing the Number of RDs

This step receives as input several complete sets of RDs, such as the one in Figure 2, which fully convey the intended concept. However, as seen in Section 3.2, these sets of RDs may contain redundant RDs. The problem of selecting the minimum number of RDs which convey the intended information is an optimization problem, which is formally expressed as follows with respect the graph in Figure 2:

$$\textbf{Minimize}\left\{\sum_{i=1}^{m} RD_i \;+ \sum_{j=m+1}^{m+p}\{RD_j\}\right\}$$

subject to: $\forall i \quad RD_i = 1 \;\; or \;\; 0$

$$\forall k \quad BEL(p_k) + \sum_{i=1}^{m} w_{i,k} RD_i \prod_{\{j | v_{j,i}=1\}} \{RD_j\} \geq T$$

The second constraint stipulates that the final belief in each proposition must be greater than a threshold T, which is determined by the attitude of the system. The final belief in a proposition p_k is composed of the previous belief in p_k plus the effect of the RDs on it. The effect of an RD RD_i on the belief in p_k is represented by the weight $w_{i,k}$, which is obtained from Step 4.4 of Propose-RDs (Section 3.2). This effect is in turn influenced by the sets of RDs which are used to convey the concepts in this RD. Each of these sets of RDs has an "all or nothing" effect. That is, if $\{RD_j\}$ is required to understand RD_i, and $\{RD_j\}$ is removed, then RD_i will have no effect on the propositions it addresses. This is represented in the $\prod$ component of the second constraint, where $v_{j,i} = 1$ if $\{RD_j\}$ is required to explain a concept in RD_i, and 0, otherwise. The values of v are obtained from Step 3 of WISHFUL (Section 3.3).

Even the easier problem of minimizing the number of RDs in the middle column of the graph, i.e., **Minimize**$\{\sum_{i=1}^{m} RD_i\}$, is NP-hard (shown by reduction to the Minimum Cover problem [Garey and Johnson, 1979]). Hence, we need a weak optimization method to minimize each alternative. To this effect, we have chosen Algorithm A* [Nilsson, 1980].

When activating A*, a node in the search graph represents the conjectured state of the hearer's beliefs, and the operators are the RDs that are applicable to a particular node. A* uses the evaluation function $f(n) = g(n)+h(n)$ for each node n in the search graph, and terminates the search at the node with the lowest value of f. The conditions for the admissibility of A* are: (1) $g^*(n) \leq g(n)$,

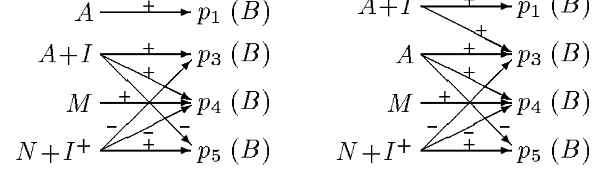

Figure 3: *Two Alternatives Generated by Propose-RDs to Convey* $\{p_1, p_3, p_5\}$

and (2) $0 \leq h(n) \leq h^*(n)$. The first condition is easily satisfied by setting $g(n)$ to the number of RDs generated up to node n. The second condition is satisfied by

h = number of *cut-sets* of the propositions still to be conveyed at node n,

where a cut-set is a group of propositions which might be conveyed by means of a single RD. That is, if the propositions to be conveyed are divided into n cut-sets, we will need at least n RDs to convey these propositions. The number of cut-sets for a set of RDs is obtained by iterating over all the RDs in the set, and for each RD, putting in the same cut-set all the propositions which are affected positively by direct or indirect inferences from this RD. Uniqueness implicatures are not considered, since by themselves they do not convey propositions, rather they perturb existing beliefs.

To illustrate the calculation of a cut-set, let us consider two of the sets of RDs which were generated by Propose-RDs to convey propositions $\{p_1, p_3, p_5\}$ (Figure 3). The main difference between these sets stems from the RDs that convey proposition p_1. In the set on the left, A affects only p_1, while in the set on the right, A+I affects both p_1 and p_3. In order to calculate the number of cut-sets in the set on the left, we first inspect A, which results in one cut-set containing p_1. Next, we inspect $A+I$, which yields a second cut-set for p_3 and p_4, meaning that these two propositions could conceivably be conveyed by means of the same RD. Note that p_5 is not in this cut-set, since it is affected negatively by A+I. Upon inspection of M, no new cut-sets are created, since p_4 already belongs to the second cut-set. Finally, when N+I^+ is inspected, a third cut-set consisting only of p_5 is generated. Like before, p_5 is not added to the second cut-set, since N+I^+ has a negative effect on p_3 and p_4. Thus, the minimum number of RDs we can possibly have in the set on the left is 3. By the same method, we reach a result of 2 for the set on the right.

After activating A*, the alternative on the left remains unchanged with 4 RDs, while the alternative on the right is reduced to 2 RDs, viz $A+I \rightarrow p_1$ and $N+I^+ \rightarrow p_5$. Note that the RD which conveys p_1 also yields a belief of *RATHER BELIEVED* with respect to p_3 (top right alternative in Figure 1). This belief is acceptable when the system is run with a daring attitude.

3.5 Selecting the Most Concise Set of RDs

When selecting the most concise among the minimized alternatives produced in the previous step, WISHFUL initially retains all the alternatives with the minimal number of RDs. For instance, since the alternative on the left in Figure 3 has 4 RDs after optimization, while that on the right has 2 RDs, the former is discarded. In order to discriminate among the remaining alternatives, those with simple RDs, e.g., A or N, are preferred to those with complex RDs, e.g., $A+I$ or $N+I^+$.

Table 3 contains the final alternative selected by WISHFUL for the example (Bracket-Simplification,

Assert + Instantiate $2(3x+5x)$	[Bracket-Simplification use-1 +/−] "In Bracket Simplification we first add or subtract the terms inside the brackets, e.g., $2(3x+5x) = 2(8x)$.
Negate + Instantiate^{+} $2(3x+5y)$	[Bracket-Simplification always apply-to Algebraic Terms] Bracket Simplification does not always apply to Algebraic Terms. For example, you cannot add the terms in brackets in $2(3x+5y)$."

Table 3: *The Set of RDs Selected by WISHFUL*

{*steps,domain*}, **daring)** which was discussed throughout this section. The possible realization of these RDs is included as an illustration.

4 Results

WISHFUL was implemented in Sun Common Lisp on a SPARCstation 2. It was run with the cautious and the daring attitude on a variety of examples and student profiles on a small animal classification and on a subset of high-school Algebra. However, owing to space limitations, we report on 8 representative trials in the Algebra domain. These trials where run with two student profiles, namely VERY GOOD and MEDIOCRE, and with a cautious and a daring attitude on the example discussed throughout this paper, and on an example where the following propositions had to be conveyed: [Distributive-Law apply-to Algebraic-Terms], [Distributive-Law apply-to Like-Terms] and [Distributive-Law apply-to Unlike-Terms]. The former example took 4 seconds of CPU time, while the latter took 1.25 seconds, with negligible variations for the different system attitudes and student profiles. The difference in the timings can be attributed to the larger number of alternatives in the longer example, and to the absence of adverse interactions between the inferences from the RDs in the shorter example.

As expected, WISHFUL produced more RDs and examples for the MEDIOCRE student than for the VERY GOOD student. In addition, WISHFUL's output showed a clear trend whereby a change from a daring attitude to a cautious attitude led to the addition of examples to the already explicit output generated for the MEDIOCRE student. This change was more dramatic for the VERY GOOD student. In this case, the daring attitude produced output such as the one in Table 3, where some of the information is implicitly conveyed, while the cautious attitude stated explicitly all the information to be conveyed, albeit with significantly less examples than for the MEDIOCRE student. This output was informally evaluated by presenting it to several lecturers and tutors in the Department of Computer Science at Monash University. There was general agreement among the interviewed staff regarding the suitability of WISHFUL's output for both types of students.

5 Conclusion

In this paper, we have offered a mechanism that achieves the seemingly incompatible objectives of generating concise discourse and addressing a hearer's inferences. These inferences have been modeled by means of inference rules, which are applied in backward reasoning mode to propose RDs that are suitable for conveying the intended information, and in forward reasoning mode to conjecture the effect of these RDs on a hearer's beliefs. These conjectures constitute an essential input to the process which optimizes the number of RDs that convey the intended information.

The discourse generated by WISHFUL is an initial explanation presented to a user. In a complete assistance system or an ITS, such an explanation is typically followed by some interaction with the user. Hence, in such a system the beliefs conjectured by WISHFUL must be verified and possibly revised after the user asks follow-up questions, or upon inspection of his/her performance.

References

[Bonarini *et al.*, 1990] A. Bonarini, E. Cappelletti, and A. Corrao. Network-based Management of Subjective Judgments: a Proposal Accepting Cyclic Dependencies. Report No. 90-067, Department of Electronics, Politecnico di Milano.

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

[Garey and Johnson, 1979] M.R. Garey, and D.S. Johnson. *Computers and Intractability, A Guide to the Theory of NP-Completeness*, W.H. Freeman and Company, San Francisco.

[Hirschberg, 1985] J.B. Hirschberg. A Theory of Scalar Implicature. Doctoral Dissertation, The Moore School of Electrical Engineering, University of Pennsylvania, Philadelphia, Pennsylvania.

[Horacek, 1991] H. Horacek. Exploiting Conversational Implicature for Generating Concise Explanations. In *Proceedings of the Meeting of the European Association of Computational Linguistics.*

[Hovy, 1988] E.H. Hovy. Planning Coherent Multisentential Text. In *Proceedings of the Twenty-Sixth Annual Meeting of the Association for Computational Linguistics*, pages 163-169, Buffalo, New York.

[Joshi *et al.*, 1984] A. Joshi, B.L. Webber, and R.M. Weischedel. Living Up to Expectations: Computing Expert Responses. In *Proceedings of the Fourth National Conference on Artificial Intelligence*, pages 169-175, Austin, Texas.

[Lascarides and Oberlander, 1992] A. Lascarides, and J. Oberlander. Abducing Temporal Discourse. In *Aspects of Automated Language Generation*, R. Dale, E. Hovy, D. Rösner, and O. Stock (Eds.), pages 167-182, Springer-Verlag, Berlin.

[Moore and Swartout, 1989] J.D. Moore, and W. Swartout. A Reactive Approach to Explanation. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 1504-1510, Detroit, Michigan.

[Nilsson, 1980] N. Nilsson. *Principles of Artificial Intelligence.* Tioga Publishing Company, Palo Alto, California.

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

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