Consulting a User Model to Address a User's Inferences during Content Planning*

Ingrid Zukerman and Richard McConachy
Department of Computer Science
Monash University
Clayton, VICTORIA 3168, AUSTRALIA
{ingrid,ricky}@bruce.cs.monash.edu.au

Abstract

Most Natural Language Generation systems developed to date assume that a user will learn only what is explicitly stated in the discourse. This assumption leads to the generation of discourse that states explicitly all the information to be conveyed, and does not address further inferences from the discourse. In this paper, we describe a student model which provides a qualitative representation of a student's beliefs and inferences, and a content planning mechanism which consults this model in order to address the above problems. Our mechanism performs inferences in backward reasoning mode to generate discourse that conveys the intended information, and in forward reasoning mode to draw conclusions from the presented information. The forward inferences enable our mechanism to address possible incorrect inferences from the discourse, and to omit information that may be easily inferred from the discourse. In addition, our mechanism improves the conciseness of the generated discourse by omitting information known by the student. The domain of our implementation is the explanation of concepts in high school algebra.

Keywords: content planning, student beliefs, inferences, backward reasoning, forward reasoning.

1 Introduction

The observation that much of what is intentionally conveyed during language use is not explicitly expressed (Grice, 1978) has been generally accepted by the Natural Language Understanding community and by researchers in Plan Recognition. Systems for discourse understanding, such as those described in (Dyer, 1982; Norvig, 1989), perform extensive inferences to understand the meaning of a piece of discourse. Similarly, plan recognition systems, such as those described in (Carberry, 1988; Litman and Allen, 1987), draw inferences from the discourse to recognize a user's intentions.

In recent times, there has been an increase in the number of Natural Language Generation (NLG) systems that take into consideration inferences which can be made from statements issued by these systems (Joshi et al., 1984; van Beek, 1987; Reiter, 1990; Zukerman, 1990a; Cawsey, 1991; Horacek, 1991; Lascarides and Oberlander, 1992). However, traditional NLG systems, e.g., (Appelt, 1982; McKeown, 1985; Paris, 1988; Moore and Swartout, 1989; Cawsey, 1990; Dale, 1990; Maybury, 1990), operate under

*This paper extends and adds detail to a paper presented by the first author at the Third International Workshop on User Modeling in Dagstuhl, Germany. The authors are indebted to the Australian Research Council for its support of this research through grant number A49030462. Sincere thanks are extended to the three anonymous reviewers for their valuable comments. The authors also thank Kate Nairn for her careful proofreading of this manuscript and helpful suggestions.

the implicit assumption that the only inferences a hearer will make from a piece of discourse are direct inferences (which reflect exactly the content of the discourse without adding to it). This assumption causes two problems:

- *Possible indirect inferences from the discourse are not addressed* – For example, a possible indirect inference from the statement "wallabies are like kangaroos" is that wallabies are the same size as kangaroos. Although this inference is incorrect, it would not be addressed by current systems.

- *The resulting text is explicit in the sense that all the beliefs which must be communicated are stated* – That is, at least one proposition is generated for each communicative goal. For instance, consider the following dialogue between a user and P.E.A., a system which gives advice about programming enhancements (Moore and Swartout, 1989):

  System (1): "You should replace `(setq x 1)` with `(setf x 1)`."
  User: "Why?"
  System (2): "I'm trying to enhance the maintainability of the program by applying transformations that enhance maintainability. `Setq-to-setf` is a transformation that enhances maintainability."

  The mechanism described in (Moore and Swartout, 1989) represents an important contribution to the field of NLG, owing to its ability to handle follow-up questions and vaguely articulated queries posed by the user. However, the explanations it generates trace the entire reasoning sequence that links the action recommended by the system with the user's goal, which in this case is enhancing the maintainability of the user's program. More concise explanations may be produced by taking advantage of the user's ability to infer some of the relationships that were explicitly stated (Horacek, 1991). This would result in text such as the following:

  System (2)': "Because we are applying transformations that enhance maintainability."
  OR
  System (2)": "Because `setq-to-setf` is a transformation that enhances maintainability."

In this paper, we present a content planning mechanism which addresses these problems. Our mechanism, which has been implemented in a system called WISHFUL, generates *Rhetorical Devices (RDs)*, such as Descriptions, Instantiations and Similes. To this effect, it consults a student model which represents three aspects of a student: (1) his/her beliefs and skills, (2) the inference rules s/he is likely to apply, and (3) his/her ability and attitude. This student model is an extension of the model discussed in (Zukerman, 1990a; Zukerman, 1990b).

The inference rules relate RDs to beliefs (Section 3.2). They are applied in two different ways during the discourse planning process: *forward reasoning* and *backward reasoning*.

- **Forward reasoning** – reasoning from the RDs to their possible effects. For instance, the application of a similarity-based inference rule to the Simile "wallabies are like kangaroos" conjectures that the user will transfer what s/he knows about kangaroos to wallabies. In order to block the transfer of features which are incorrect with respect to wallabies, such as size, a disclaimer, such as "but smaller," must be added. This reasoning mechanism was used in (Zukerman, 1990a) for the generation of Contradictions and Revisions to possible inferences drawn by a student.

- **Backward reasoning** – reasoning from the goal to be accomplished 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 Example, or a combination of these RDs. This reasoning mechanism has been widely used in NLG systems, e.g., (Appelt, 1982; Hovy, 1988; Moore and Swartout, 1989; Cawsey, 1990; Maybury, 1990). Some of these systems, e.g., (Moore and Swartout, 1989; Cawsey, 1990), have encoded particular inferences into discourse planning operators. However, these operators are applicable only in one direction, namely backward reasoning. In addition, this encoding does not represent explicitly the inferential process that allows a user to deduce a belief from an RD. For instance, the following operator (Moore, 1989) supports the generation of an exhaustive set of examples to convey a concept, but it does not indicate why the user will infer the concept in question from these examples.

```
NAME: describe-by-example
EFFECT: (BEL ?hearer (CONCEPT ?concept))
CONSTRAINTS: (AND (ISA ?concept OBJECT)
                  (IMMEDIATE-SUBCLASS ?example ?concept))
NUCLEUS: (FORALL ?example
                 (ELABORATE-CONCEPT-EXAMPLE ?concept ?example))
SATELLITES: nil
```

Our content planner also follows Grice's Maxim of Quantity (Grice, 1975) in that it omits information known by the student. To this effect, it consults our model of the student's beliefs and skills. This feature is particularly useful in situations such as those described in (Sleeman, 1984), where a student knows most of the steps in a procedure, and needs to be instructed only with respect to a few of them.

In the next section, we discuss previous research that focuses on addressing a user's inferences during discourse planning. In Section 3, we describe our student model. In the remainder of the paper, we describe the tasks performed by the content planner, and the contribution of the student model to each of these tasks.

2 Related Research

The research reported in (Joshi et al., 1984; van Beek, 1987; Zukerman 1990a) considers the addition of information to planned discourse to prevent or weaken a user's erroneous inferences from this discourse.

Joshi et al. (1984) and van Beek (1987) characterize situations where explanations must be added to expert responses to a user's queries in order to block a user's erroneous inferences from these responses. Zukerman (1990a) adds Contradictions and Revisions to planned propositions based on the conjectured effect of the user's inferences from these propositions on his/her beliefs.

The research described in (Horacek, 1991; Lascarides and Oberlander, 1992) considers the omission of information that may be inferred by the user from planned discourse. Horacek (1991) omits domain-related information from the explanation of the solution of constraint satisfaction problems if this information may be inferred by the user from the explanation (possibly in combination with the user's domain knowledge). Lascarides and Oberlander (1992) remove temporal information that may be easily inferred from the manner in which discourse is presented.

Finally, Cawsey (1991) takes into consideration inferences which result from the inheritance of attributes in hierarchical domains in order to convey the attributes of objects by means of Similes and Instantiations rather than Descriptions.

3 The Student Model

Our student model is composed of three main parts: (1) representation of a student's beliefs and skills, (2) representation of his/her possible inferences, and (3) representation of the student's ability and attitude.

3.1 Representation of a Student's Beliefs and Skills

The representation of beliefs and skills in our student model is based on the representation described in (Zukerman, 1990a; Zukerman, 1990b). The beliefs and skills in our model pertain to technical information, which is characterized by the presence of procedures which achieve certain goals when applied to particular objects.

Our model distinguishes between two aspects of relations which involve actions: (1) the student's belief in the correctness of these relations, and (2) the student's skill with respect to the actions specified in these relations. This distinction allows us to represent situations where a student believes a proposition to be correct, even though s/he lacks the skill to perform the action mentioned in this proposition. For example, a student may believe that [+ apply-to Like-Terms], even though s/he may not know how to perform addition of Like Terms.

In order to represent propositions whose correctness depends on the truth of other propositions, we require an explicit representation of context. For example, Factorization yields a product of factors when applied to decomposable expressions, i.e., expressions of the form $amx^2+(bm+an)x+bn$. However, when Factorization is applied to non-decomposable expressions, it yields an expression that is not a product of factors. We use the following notation to represent the circumstances under which a relation holds:

[concept$_1$ relation concept$_2$ context]. The meaning of this representation is that the *relation* between *concept$_1$* and *concept$_2$* holds in a particular *context*, where the *context* is either the global context or an arbitrary sequence of nested relationships. Such a sequence provides a uniform representation for chains of conditions. In our example, the relation [Factorization has-goal Product-of-Factors (Factorization apply-to Decomposable-Expressions)] means that Factorization will produce a product of factors when applied to expressions of the form $amx^2 + (bm + an)x + bn$.

Our model is implemented by means of a network whose nodes represent concepts, and whose links represent relationships between concepts. Figure 1, adapted from (Zukerman, 1990a), depicts a network that represents part of the knowledge of a student who has been taught the steps of Bracket Simplification, and told that Bracket Simplification applies to Numbers. In addition, this network represents the assumption that the student has inferred correctly that Bracket Simplification also applies to Like Terms, and inferred incorrectly that Bracket Simplification applies to Algebraic Terms and Unlike Terms.

The network in Figure 1 contains the objects Numbers, Algebraic-Terms, Like-Terms and Unlike-Terms (represented by ovals); the procedures Bracket-Simplification, $\times$ and $+/-$ (represented by rectangles); and the goal state Brackets-Eliminated (represented by an oval). The links in this network are labelled with the predicates apply-to, use-i, has-goal, isa and similar. The relation [P use-i Q] means that Q is the ith step of P. For instance, the use-1 link of Bracket Simplification indicates that the first step of Bracket Simplification is addition or subtraction ($+/-$), and the use-2 link indicates that the second step is multiplication. Thus, given an expression composed of Like Terms, such as $2(4x + 5x)$, the Bracket Simplification procedure first adds the terms in the brackets, yielding $2(9x)$, and then multiplies the result in brackets by the factor outside the brackets, yielding $18x$ (an expression without brackets). Contextual information is represented by attaching a qualifier to the predicate which labels a link. For instance, the qualifier (BrS apply-to N) attached to one of the has-goal links between Bracket-Simplification and Brackets-Eliminated indicates that Bracket Simplification achieves the goal of eliminating brackets when applied to Numbers. Similarly, the qualifiers of the remaining has-goal links, i.e., (BrS apply-to AT), (BrS apply-to LT) and (BrS apply-to UT), indicate that Bracket Simplification achieves the goal of bracket elimination when applied to Algebraic Terms, Like Terms and Unlike Terms, respectively. Both the links and the nodes are labelled according to the manner in which they were acquired by the student, i.e., **I**nferred (by means of an inference rule), **T**old (by the system) or **K**nown Previously. In Figure 1, Inferred links have normal thickness, while links that are Told and links that are Known Previously appear in boldface.

The information in the student model is represented at a level of detail which is consistent with the level of expertise required to learn the subject at hand. That is, well-known concepts, such as $\times$ and $+/-$, are represented by singleton nodes, while relatively new concepts, such as Bracket-Simplification, are broken down into their components. This level of detail is initially determined by the designers of the

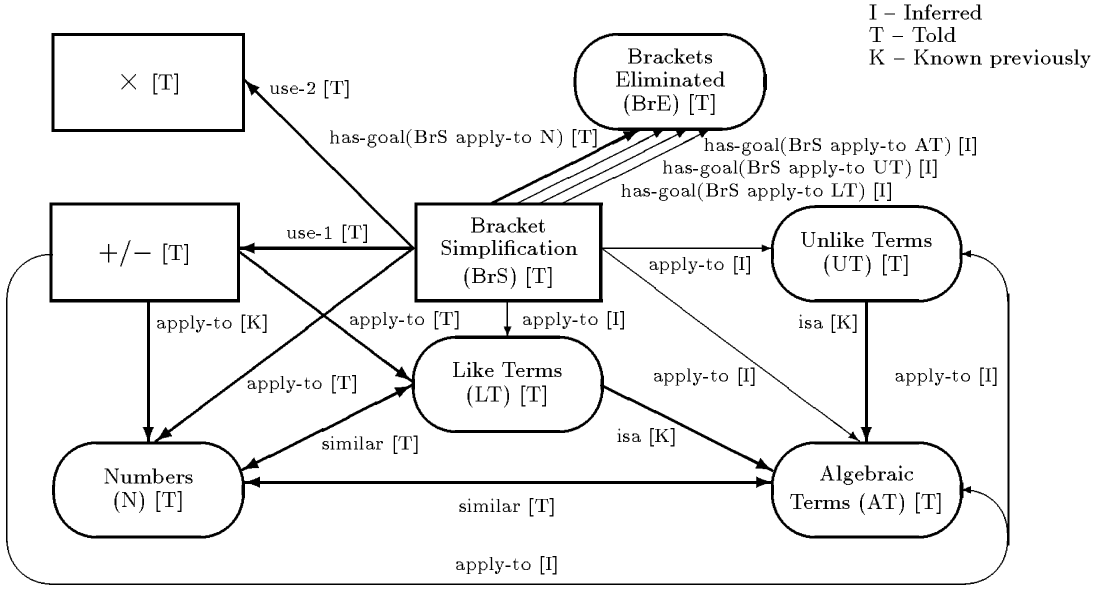

Figure 1: Network Model of a Student's Beliefs

network. The development of a mechanism which automatically adjusts the level of detail of a network to reflect the progress of a student is the subject of future research.

3.1.1 Representation of the Strength of a Student's Beliefs

Each of the links and nodes in the student model is assigned a value which represents the strength of a student's conjectured belief in the information in question. In the current implementation, this value is represented by means of the following *qualitative belief states* (Driankov, 1986; Bonarini et al., 1990): {*DISBELIEVED (D), RATHER DISBELIEVED (RD), UNKNOWN (U), CONTRADICTORY (C), RATHER BELIEVED (RB), BELIEVED (B)*}. A qualitative representation is justified by the type of assessment of a student's beliefs and skills that is useful for discourse planning, i.e., a broad assessment, rather than a pinpoint numerical assessment. For example, students with widely different capabilities and attitudes will usually receive explanations which differ in their approach and level of detail (Paris, 1988), while students with similar capabilities are likely to receive similar explanations. This representation also avoids other problems inherent in numerical methods (Clark, 1990) without forcing the abandonment of some of the useful underpinnings of such methods.

Driankov's qualitative belief states are composed of two related measures:

- **Support** $s(A)$ – which gives the positive evidence for a proposition A.
- **Plausibility** $p(A)$ – which is defined as $p(A) = 1 - s(\neg A)$, i.e., the absolute certainty of A minus the support for the negation of A. The plausibility of a proposition A reflects the user's certainty with respect to the evidence for A.

The maintenance of separate measurements for support and plausibility has clear semantics, thereby avoiding the ambiguity which results from having a single number associated with a belief. Both of these measures are represented in the system by fuzzy numbers in the interval $[0-1]$. It is envisioned that for each student or each type of student, a teacher will input the initial values of these measures for all the concepts and assertions that are relevant to the material being taught. After interacting with the student, the system will update these values using operators defined in (Bonarini et al., 1990).

Support for a proposition may come from many different pieces of evidence. The level of support resulting from several individual pieces of evidence depends on (1) the number of pieces of evidence, and (2) the support that each piece of evidence lends to the proposition. The plausibility of a proposition qualifies the support for this proposition. The plausibility of A will be high when there is little or no evidence against it, and low when the support for $\neg A$ is high. If there are large quantities of unreliable evidence for a proposition A, the aggregation of this evidence will result in A having a high level of support. However, since the evidence is unreliable, thereby having a low plausibility, A will also have a low plausibility.

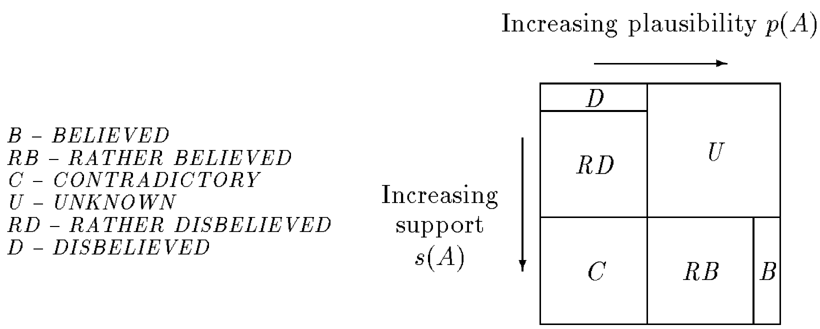

Figure 2: The Effect of Support and Plausibility on the Belief State

The measures of support and plausibility are combined as shown in Figure 2 (Bonarini et al., 1990) to give the six possible belief states. For example, if we have significant evidence for both of the propositions A and $\neg A$, then both $s(A)$ and $s(\neg A)$ will be large, but both $p(A)$ and $p(\neg A)$ will be small. In this case, from Figure 2 we see that both the belief in A and the belief in $\neg A$ are represented as *CONTRADICTORY*. This agrees with what most people would think when faced with evidence in favour of two mutually exclusive events.

In our research, we have used the measures of belief and the combination rules as defined in (Bonarini et al., 1990). However, it is possible that in borderline cases, these combination rules will lead to results that do not match human intuition. For example, consider a proposition X with values for $s(X)$ and $p(X)$ that place the belief in X at the junction of the lines intersecting between the belief states *RD, U, RB* and *C* in Figure 2. Now, if new evidence for X is found, $s(X)$ will increase, and the new belief state will be on the line separating the states *C* and *RB* somewhere below the original belief state. This means that although the support for X has increased, the belief state will remain essentially unchanged. This problem can be solved by slightly tilting the two long intersecting lines in Figure 2 so that they both have a positive gradient. In our example, this will result in the belief in X shifting to RB as the support for X increases. The consideration of the effect of these modifications on the original theory is the subject of future research.

3.2 Representation of a Student's Inferences

Our mechanism takes into consideration three types of inferences: (1) *direct inferences*, (2) *indirect inferences*, and (3) *uniqueness implicatures*. These inferences are represented by means of inference rules which have the following general format:

$$\textit{Inference(RD,[beliefs])} \xrightarrow{L_{rule}, C_{rule}} \textit{Belief} \text{ and/or } \textit{Skill},$$

where L_{rule} is the likelihood that a student will retain an inference drawn by a particular inference rule, and C_{rule} is the confidence the student has in the rule. That is, a rule infers a *Belief* and/or a *Skill* from an RD possibly in combination with beliefs already held by the student. This belief (or skill) is retained with likelihood L_{rule} and rejected otherwise. Further, the confidence in the retained belief (or skill) is adjusted by the factor C_{rule}. For instance, the rule Generalize(Assert($P(Z)$),[Z inst Z']) $\xrightarrow{L_{G+}, C_{G+}}$ Belief($P(Z')$) states that given an Assertion of proposition $P(Z)$ and a student's belief that Z is an instance of Z', the student will conclude with probability L_{G+} that $P(Z')$ is true. Further, the student's belief in $P(Z')$ is a function of C_{G+} and of his/her belief in the proposition [Z inst Z']. To illustrate the application of this rule, consider the Assertion of the proposition [Bracket-Simplification use-1 $+/-$] accompanied by an Instantiation with respect to the expression $2(4x + 3x)$. This RD results in text such as "In bracket simplification we first add or subtract the terms inside the brackets, e.g., $2(4x + 3x) = 2(7x)$." From this RD and the belief that [$2(4x + 3x)$ inst Like-Terms], the generalization inference rule will yield the belief that [Bracket-Simplification apply-to Like-Terms].

In this manner, our inference rules allow our system to conjecture the effect of an RD on a student's beliefs, and act accordingly, i.e., omit information that may be inferred from this RD, and add information that addresses incorrect inferences from the RD. In the current version of WISHFUL, the beliefs represented in the student model may be updated only by means of the inferences postulated by the inference rules. Clearly, such a model may eventually diverge from the real status of the student beliefs. Hence, in a fully interactive system, the beliefs postulated by the student model should be validated against the student's performance.

3.2.1 Direct Inferences

A direct inference from an RD yields an understanding of the information in this RD. In our model, this means that when proposition P is asserted, a belief in P is generated, and when P is negated, a belief in $\neg P$ is generated. The Abstract-Understand inference rule assesses the likelihood that a student will understand an Assertion or a Negation by means of a direct inference and the confidence the student will place in the resulting belief and/or skill (Table 1[1]). Both of these factors are affected by (1) the abstractness and complexity of the proposition to be conveyed, which depend on the complexity of the nodes that participate in the proposition (Section 3.1); and (2) the addressee's ability to understand abstract and/or complex information. For instance, a capable student is likely to understand an abstract statement, while a mediocre student will be lost if a more concrete explanation does not accompany the abstract statement.

Abstract-Understand (AU) –
Abstract-Understand(Assert(P)) $\xrightarrow{L_{AU+},C_{AU+}}$ Belief(P) and/or Skill(P)
Abstract-Understand(Negate(P)) $\xrightarrow{L_{AU-},C_{AU-}}$ Belief($\neg P$)
Generalization (G) –
Generalize(Assert($P(Z)$),[beliefs]) $\xrightarrow{L_{G+},C_{G+}}$ Belief($P(Z')$) and/or Skill($P(Z')$)
Generalize(Negate($P(Z)$),[beliefs]) $\xrightarrow{L_{G-},C_{G-}}$ Belief($\neg P(Z')$),
where Z' is a super-class of Z.
Specialization (S) –
Specialize(Assert($P(Z)$),[beliefs]) $\xrightarrow{L_{S+},G_{S+}}$ Belief($P(Z')$) and/or Skill($P(Z')$)
Specialize(Negate($P(Z)$),[beliefs]) $\xrightarrow{L_{S-},G_{S-}}$ Belief($\neg P(Z')$),
where Z' is a subclass of Z.
Uniqueness Implicature (UI) –
Uniqueness-Implicate(Assert($P(S,O)$),[beliefs]) $\xrightarrow{L_{UI},C_{UI}}$ $\forall x \neq O$ Belief($\neg P(S,x)$)

Table 1: Sample Inference Rules Used in WISHFUL

3.2.2 Indirect Inferences

Indirect inferences draw conclusions that are removed from what was said by one or more inference steps. The indirect inference rules considered at present in our model are based on the rules described in (Zukerman, 1990a; Zukerman, 1990b), namely: generalization, specialization (Table 1), similarity and applicability. The first three rules reflect student behaviour observed in (Matz, 1982). The generalization rule was also postulated in (van Lehn, 1983; Sleeman, 1984). The similarity rule transfers the attributes of a source concept to a target concept, e.g., from kangaroos to wallabies. The applicability rule is a simple deductive reasoning rule. It states that if the first set of steps of a procedure is applicable to an object of a certain type, then the entire procedure is applicable to this object. For example, since addition and subtraction are applicable to Numbers, and the first step of Bracket Simplification is addition or subtraction, this rule allows us to conclude that Bracket Simplification applies to Numbers.

The likelihood that an indirect inference will be retained by a student is affected by the following

[1]The $+$ and $-$ subscripts represent inferences from Assertions and Negations respectively.

factors: (1) the student's ability and attitude (Section 3.3); (2) the correctness of the inference; and (3) the soundness of the rule that yields the inference. The first two factors are used to model the behaviour observed by Sleeman whereby good students retain more correct conclusions than mediocre students (Sleeman, 1984). For instance, given the Instantiation $(x-3)(x-4)=0 \Rightarrow x-3=0$ or $x-4=0$, a student may perform the wrong generalization $(x-A)(x-B)=K \Rightarrow x-A=K$ or $x-B=K$ (Matz, 1982). According to Sleeman, both a good and a mediocre student may perform this mis-generalization. However, the good student will be more critical of the conclusion than the mediocre student, and will usually discard the incorrect conclusion. The mediocre student, on the other hand, is more likely to retain an incorrect inference obtained in this manner. The third factor is necessary in order to model the amount of faith different types of students place in different types of inference rules. For example, a good student is more likely to retain the conclusions drawn by sound inference rules than the conclusions drawn by unsound rules. On the other hand, a mediocre student may be unable to discriminate between sound and unsound inference rules, making him/her equally likely to accept the conclusions drawn by both types of rules.

The factors that affect a student's confidence in a conclusion from an indirect inference are: (1) the soundness of the inference rule that yields the conclusion, (2) the student's ability and attitude, (3) the strength of the student's beliefs which participate in the indirect inference in question, and (4) the strength of the existing belief in the conclusion. For instance, given the Assertion "kangaroos hop," a student's confidence in the proposition "wallabies hop" depends on the student's knowledge about the relationship between wallabies and kangaroos and on his/her confidence in the similarity inference rule, which in turn depends on the soundness of this rule and on the student's ability and attitude. Thus, the first two of the above factors determine the confidence factor (C) that the system assigns to each of the different types of inference rules for each type of student. For example, a good student may be rather cautious with respect to a conclusion drawn by an unsound inference rule, while a mediocre student may believe this conclusion more strongly (Section 3.3). If a (retained) conclusion affects a belief currently held by a student, then the plausibility and support of the conclusion will be combined with the plausibility and support of the student's existing belief as described in (Bonarini et al., 1990).

Indirect inferences are categorized into three types based on their soundness: (1) *sound*, (2) *acceptable*, and (3) *unacceptable*.

- **Sound inferences** – inferences which are logically sound, such as a specialization from a positive statement or a generalization from a negative statement. For example, the Assertion "Marsupials are indigenous of Australia" specializes to "Kangaroos are indigenous of Australia," and the Negation "Bracket Simplification does not apply to Unlike Terms" generalizes to "Bracket Simplification does not always apply to Algebraic Terms."

- **Acceptable inferences** – common-sense inferences whose results hold most of the time, e.g., a generalization from a positive instance to a class or a specialization from a negative statement.

- **Unacceptable inferences** – inferences whose results hold only sometimes, and hence should not be sanctioned, e.g., the transfer of features between two items with superficial similarities, without the transfer being specifically suggested by means of a Simile.

Our characterization of acceptable and unacceptable inferences is incomplete in the sense that there are factors other than frequency that affect the acceptability of an unsound inference. For example, in the case of a generalization from a positive statement, the typicality of the subclass or the instance from which the generalization is made affects the correctness of the resulting inference, e.g., since $3x + 5y$ is more typical of Algebraic Terms than $3x+5x$, a generalization from $3x+5y$ to all Algebraic Terms is more likely to be correct than a generalization from $3x + 5x$. Similarly, similarity-based inferences between items that are in close proximity in a concept hierarchy are more likely to be correct than similarity-based inferences between items that are far apart in a concept hierarchy. For instance, a similarity-based inference from kangaroos to wallabies is more likely to yield correct conclusions than a similarity-based inference from canaries to penguins. The consideration of the effect of these factors on our inference rules is left for future research.

3.2.3 Uniqueness Implicatures

Given an asserted proposition $P(S, O)$, a uniqueness implicature licenses the inference that O is the *only* instance of x for which $P(S, x)$ is true. For example, upon hearing the statement "Joe has one leg," most people will infer that Joe has one leg *only* (Hirschberg, 1985).

Several researchers have addressed context dependent implicatures by means of rules or operators which embody the characteristics of a situation. Hirschberg (1985) provided rules that license scalar implicatures, which are based on an ordering between the entities affected by a piece of discourse. Horacek (1991) used rules which license implicatures from the information in a user's query and in the reply to this query in order to omit information from this reply. Green and Carberry (1992) used discourse plan operators which are similar to Hirschberg's licensing rules in order to interpret and generate replies that involve scalar implicatures. Finally, Reiter (1990) embodied implicatures from lexical items in rules for lexical selection.

A uniqueness implicature is a type of scalar implicature which is particularly relevant to knowledge acquisition settings. This is because in these settings, the aim is to extend a student's knowledge, but infelicitous wording could lead to a false uniqueness implicature which would replace rather than extend correct beliefs held by a student. For instance, if a speaker says "Bracket Simplification applies to Like Terms," a uniqueness implicature will license the inference that Bracket Simplification applies

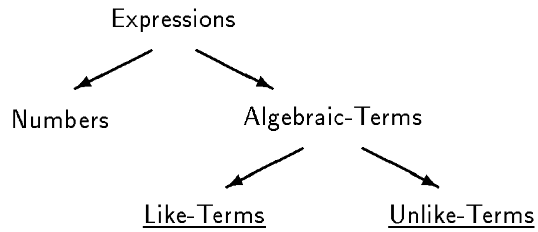

Figure 3: Partial Algebraic Concept Hierarchy

only to Like Terms. Now, if the student believes that Bracket Simplification applies to Numbers, the uniqueness implicature will cause a conflict with this belief. In order to address this conflict, an RD which acknowledges the correctness of the student's belief must be included in the discourse, e.g., "Bracket Simplification applies to Like Terms, *as well as to Numbers*." We call this type of RD *Mention*.

Given the RD **Assert**$(P(S, O))$, a false uniqueness implicature is anticipated by conducting a top-down traversal of the hierarchical relations in the network representing the student's beliefs, and retrieving the highest concepts in the hierarchy, $\{C_1, \ldots, C_n\}$, for which the student correctly believes that $P(S, C_i)$ holds. These concepts are then referred to when the uniqueness implicature is addressed. For example, consider the asserted proposition [Distributive-Law apply-to Numbers] and the concept hierarchy in Figure 3 (the underlined concepts are those for which the student correctly believes in the applicability of Distributive Law). After traversing the concept hierarchy, the concepts {Like-Terms, Unlike-Terms} are returned. Note that although it is true that [Distributive-Law apply-to Algebraic-Terms], Algebraic-Terms is not returned, since the student does not believe in the correctness of this proposition. The resulting RD is **Mention** [Distributive-Law apply-to {Like-Terms, Unlike-Terms}], which may be realized as "*In addition to Like Terms and Unlike Terms*, Distributive Law applies to Numbers."

Uniqueness implicatures differ from the other types of inferences in that they do not yield lasting beliefs in memory, rather they cause conflicts with existing beliefs. For example, if a student has no previous beliefs with respect to a proposition P, a uniqueness implicature from the RD **Assert**$(P(S, O_1))$ does *not* result in the permanent belief that $P(S, x)$ is false for all objects x. Hence, if $P(S, O_2)$ is asserted later on, a false uniqueness implicature will contradict only the student's belief in $P(S, O_1)$.

3.3 Modeling a Student's Ability and Attitude

At present, the system has five predefined profiles which represent a student's ability: *EXCELLENT, GOOD, AVERAGE, MEDIOCRE* and *BAD*[2]. When we apply the inference rules discussed in Section 3.2 in the context of these profiles, the results emulate behaviour observed in (Sleeman, 1984), whereby good students are discerning about the inferences they make and retain more correct conclusions than incorrect ones, while the opposite happens for mediocre students. Informal observations of students indicate that

[2] The names associated with the different profiles are not value judgments of individual students, rather these categories represent different levels of correctness and conviction of a student's presumed beliefs and inferences.

good students have more faith in their indirect reasoning skills than mediocre students. That is, stronger students are more certain of the results of their indirect inferences than weaker students. This observation was also incorporated in the student profiles.

The student profiles are implemented as files with entries for the likelihoods and confidence factors of the different inference rules. These likelihoods and confidence factors were set to reflect common behaviour patterns observed in students, and were adjusted empirically based on the generated output. Table 2 illustrates the likelihoods and confidence factors of the inference rules used by an *EXCELLENT* student. As one would expect from an *EXCELLENT* student, the likelihood of retaining an inference is much higher when it is correct than when it is incorrect. The likelihoods for weaker students reflect a different behaviour, where the proportion of correct to incorrect conclusions that are retained is less favorable than for good students. L_{G-}, L_{S+} and L_D always have a value of 0 for incorrect inferences, since generalizations of Negations, specializations of Assertions and deductive inferences are sound, and hence can never be incorrect. L_{AU} has a value for absent inferences rather than for incorrect ones, because the system models lack of learning from a direct inference, rather than mislearning. That is, according to our model, a student may or may not understand an Assertion or a Negation to the required level, but s/he will not learn the wrong thing. We assume that *EXCELLENT* students are confident with respect to the inferences they draw by means of sound and acceptable inference rules. However, they are less confident with respect to unsound inferences. Hence, the confidence factor is 1 for all the inference rules except for the Similarity-based inferences. As stated before, weaker students may have different confidence factors associated with the different types of inference rules, e.g., they may have more confidence in unsound rules than the stronger students. Finally, L_{UI} has a value of 1, because in the current implementation the student is assumed to draw uniqueness implicatures from all Assertions. Changing this assumption so that only certain false uniqueness implicatures are anticipated and addressed is a current research problem.

Soundness of an Inference	**Inference Rule**	**Inference is**				
		Correct		**Incorrect**		**Absent**
		L	C	L	C	L
	Abstract-Understand (AU)	0.9	1			0.1
Sound	Generalization from Negation ($G-$)	0.95	1	0	0	
	Specialization from Assertion ($S+$)	0.95	1	0	0	
	Deduction (D)	0.95	1	0	0	
Acceptable	Generalization from Assertion ($G+$)	0.9	1	0.2	1	
	Specialization from Negation ($S-$)	0.9	1	0.2	1	
Unacceptable	Similarity (Sim)	0.75	0.5	0.25	0.5	
	Uniqueness Implicature (UI)	1	1	1	1	

Table 2: Sample L and C Values for an *EXCELLENT* Student

The five student types are further refined using the attitude modifiers *ABSTRACTLY-INCLINED* and *CONFIDENT* and their negations (Table 3). For example, an *EXCELLENT* student is usually *ABSTRACTLY-INCLINED* and *CONFIDENT* about his/her conclusions. Changing this student's attitude to *NOT-ABSTRACTLY-INCLINED* has the effect of making our student no better at understanding complicated abstract statements than an *AVERAGE* student. However, a student of this type will still be more capable than an *AVERAGE* student at other sorts of reasoning, such as reasoning involving indirect inferences. The modifier *NOT-ABSTRACTLY-INCLINED* is implemented for the *EXCELLENT* student as follows: the L_{AU} and C_{AU} factors for correct inferences are reduced to reflect a lower likelihood of understanding an abstract statement and a lower level of confidence in the understood information, respectively. In addition, the L_{AU} factor for absent inferences is raised to a value similar to that of an *AVERAGE* student.

Modifier	**Main Effect**
ABSTRACTLY-INCLINED	Increase the student's ability to understand abstract explanations. (Increase L_{AU} and C_{AU} for correct inferences, and decrease L_{AU} for absent inferences.)
NOT-ABSTRACTLY-INCLINED	Opposite to the above.
CONFIDENT	Increase the student's conviction in any inferences s/he makes. (Increase the C values of all the inference rules.)
TIMID	Opposite to the above.

Table 3: Student Attitude Modifiers

The predefined student classes express conveniently the characteristics of common stereotypical students. In addition, the modification of student types by means of attitudes supports the maintenance of many different student profiles. There are several ways to choose an initial profile for a student. One way consists of allowing a teacher to determine which profile fits best a particular student. In a fully operational interactive system, the system could deduce a profile by querying the student, or alternatively it could simply start with an *AVERAGE* profile and update it based on the student's performance. New student profiles and modifiers can be easily added to the system, since they are merely input files which contain values for the various likelihoods and confidence factors used in the inference rules.

The clear separation between the profile of the student's ability and attitude and the model of the concepts and relations known by the student enables us to model a wide range of students. At present, the student's type and modifiers are kept constant, while the network representing the student's factual knowledge base is allowed to grow as the student learns new information. This represents a student who is able to learn new facts, but whose reasoning ability does not improve as a result of his/her new factual knowledge. An interesting avenue of future research involves simulating a student whose modifiers and type change gradually as his/her knowledge base grows, thereby modeling a student whose deductive powers improve as his/her knowledge increases.

3.4 Comparison with Our Previous Work

As stated above, the student model presented in this paper is based on the model discussed in (Zukerman, 1990a; Zukerman, 1990b). However, our current model differs from the previous model in the following aspects:

- **Representation of a student's ability and attitude** – This aspect, which is required in order to express the differences between various types of students, did not exist in our previous model.

- **Representation of a student's beliefs and skills** – The parts of our representation that differ from the previous representation are: (1) the distinction between beliefs and skills, (2) the representation of contextual information, and (3) the representation of the strength of a student's conjectured beliefs. As stated in Section 3.1, the distinction between beliefs and skills is necessary in order to represent these two facets of a student's knowledge. The representation of contextual information is necessary in order to represent propositions whose correctness depends on the truth of other propositions (Section 3.1). Finally, our previous representation of the strength of a student's belief was similar to the representation used in MYCIN (Buchanan and Shortliffe, 1985), while our current representation is based on the qualitative belief states developed by Driankov (1986) and expanded by Bonarini et al. (1990). The reasons for the shift in representation are twofold: (1) the inadequacy of MYCIN's Certainty Factors for combining several pieces of evidence (Buchanan and Shortliffe, 1985), and (2) the need to make a broad assessment of a student's beliefs and skills.

- **Representation of a student's inferences** – The explicit representation of direct inferences and the inclusion of uniqueness implicatures constitutes an expansion of our previous representation. In addition, our inference rules differ from those described in our previous model in their domain. The current rules draw inferences from RDs possibly in combination with beliefs already held by the student, while the rules used in the earlier research draw inferences from already acquired beliefs only. This change in domain is necessary since drawing inferences only from beliefs and not from RDs presupposes the beliefs that will be inferred from an RD.

4 Operation of the Content Planner

Our content planner receives as input a **concept** to be conveyed to a student (e.g., Distributive Law), a list of **aspects** that must be conveyed about this concept (e.g., operation and domain), and a **communicative goal**, which states the degree to which these aspects must be known by the student (e.g., well known). This type of information can be provided by an Intelligent Tutoring System (ITS), but in our system it is hand-coded. The output of the content planner is a list of RDs, where each RD is composed of a rhetorical action, such as Assert or Instantiate, applied to a proposition (Section 4.2).

In order to convey the intended aspects of a concept, our mechanism first determines the information to be presented, and then proposes RDs to convey this information. However, it is possible that the student does not understand the concepts mentioned in a particular RD well enough to understand this RD. Therefore, the generation process is repeated with new communicative goals and aspects with respect to the concepts mentioned in the proposed RDs, in order to present information about these concepts if necessary.

The block diagram in Figure 4 illustrates the tasks performed by the content planning process. In the following sub-sections, we discuss these tasks, with particular reference to the following sample input: (Bracket-Simplification, **KNOW**, {*domain,operation*}). That is, the communicative goal is for the student to know the domain and operation of the Bracket Simplification procedure. Other discourse planning tasks, such as organizing the generated RDs and generating referring expressions, are constrained by the outcome of the content planning process. However, they are not an integral part of this process. Rather, they are separate tasks which deserve independent consideration (Suthers, 1991).

WISHFUL: concept, communicative-goal, aspects

1. Decide which information to present
(Map Filter)

↓

2. Propose Rhetorical Devices
(Address inferences)

↓

3. For each concept in the generated messages:
WISHFUL(concept, communicative-goal', aspects')

Figure 4: Tasks Performed by the Content Planner

Our content planner works in the paradigm of first deciding what to say, and then determining how to say it, where the latter decision is left to a separate text realization component. An alternative paradigm consists of interleaving these decisions, since after deciding what to say, it may not be possible to generate a legal text that actually conveys the intended information (Appelt, 1982; Meteer, 1991). The first paradigm was adopted for our content planner because interleaving content planning with discourse realization would obscure the issues addressed in this research.

4.1 Deciding which Information to Present

The content planner stresses the presentation of information that the student does not know and information that addresses misconceptions held by the student. In order to generate such a discourse, our system consults our model of the student's beliefs and skills. The following procedures are applied to perform this task: *Mapping* and *Filtering*.

4.1.1 Mapping

The mapping procedure expresses the aspects to be conveyed with respect to a concept in a manner which is compatible with the student model. For instance, in the current implementation, aspects such as domain and operation are mapped into predicates such as apply-to and use-i for $i = 1, ..., n$ (n is the number of steps in a procedure), respectively. According to this mapping, the above input to the system (Bracket-Simplification,**KNOW**,{*domain*,*operation*}) is mapped into the following propositions:

Aspect	Domain Predicate
domain	[Bracket-Simplification apply-to Like-Terms] [Bracket-Simplification apply-to Numbers]
operation	[Bracket-Simplification use-1 $+/-$] [Bracket-Simplification use-2 $\times$]

Table 4: Propositions Relevant to the Sample Aspects

As seen in this example, the mapping process is often trivial. However, it is required for theoretical clarity, as it distinguishes between high level didactic decisions regarding the aspects to be conveyed about a concept and the representation of these aspects in the student model.

4.1.2 Filtering

In this step, the system removes the propositions that are already known by the student from the list of propositions to be conveyed, and adds to this list propositions which correct information that is wrongly believed by the student with respect to the given aspects. This process expands on the process discussed in (Moore and Paris, 1992), where a user model is consulted in order to omit from the discourse information that the user is presumed to know. Propositions that are weakly believed by the student are presented, but they must be prefixed with a Meta Comment which credits the student with the belief in question (Zukerman, 1991), e.g., "*As you probably know*, Bracket Simplification applies to Numbers." This process is particularly useful when a student's knowledge is lacking with respect to a few items of information only.

To illustrate the filtering process, consider a situation where the student believes the following with respect to Bracket Simplification:

[Bracket-Simplification apply-to Algebraic-Terms] [Bracket-Simplification apply-to Numbers] [Bracket-Simplification use-2 $\times$]

In this case, the propositions [Bracket-Simplification apply-to Numbers] and [Bracket-Simplification use-2 $\times$] are filtered out from the list of propositions to be conveyed. In addition, the negation of the proposition [Bracket-Simplification apply-to Algebraic-Terms] is added to the list of propositions to be conveyed, since it is incorrectly believed by the student. This process results in the list of propositions in Table 5.

Aspect	Domain Predicate
domain	[Bracket-Simplification apply-to Like-Terms] [Bracket-Simplification ¬(always) apply-to Algebraic-Terms]
operation	[Bracket-Simplification use-1 +/−]

Table 5: Propositions to be Conveyed

4.2 Proposing Rhetorical Devices

In this step, the content planner proposes RDs to convey the set of propositions produced in the previous step. To this effect, it takes into consideration inferences a user is likely to perform from these RDs. Our procedure is based on the tenet that while processing a piece of discourse in an interactive setting, a user will draw immediate inferences from the discourse, but will perform further reaching inferences only after the entire discourse has been processed. In order to address these immediate inferences, our procedure draws one round of inferences from a proposed RD. For instance, given the RD **Assert**[Bracket-Simplification apply-to Like-Terms], the generalization rule produces the incorrect inference [Bracket-Simplification apply-to Algebraic-Terms], and the similarity rule yields the correct inference [Bracket-Simplification apply-to Numbers]. The content planner then omits from the planned discourse the correct inferences the user is likely to make, and generates discourse which addresses the incorrect inferences. This process is carried out by the procedure *Propose-RDs*, which produces a list of RDs, i.e., rhetorical actions applied to propositions.

Propose-RDs(list-of-propositions-to-be-conveyed, aspect)

1. If *aspect* = *nil* or there are no more propositions to be conveyed that pertain to *aspect*, then

   Select an aspect to be conveyed, and assign it to *aspect*.

2. Select a proposition which pertains to *aspect*.

3. Apply inference rules in *backward reasoning* mode in order to propose a set of RDs which convey this proposition.

   (Each RD in this set constitutes a different alternative for conveying the proposition in question.)

4. For each alternative RD in the set of RDs, apply inference rules in *forward reasoning* mode in order to draw the inferences that can be made from this RD.

   (a) Update the list of propositions to be conveyed as follows:

      i. If an inference is correct and it corresponds to one of the propositions to be conveyed, then

A. If the inference is strong enough then mark the proposition as *deleted from the list of propositions to be conveyed*, i.e., it no longer has to be said.

(An inference is strong enough if it satisfies the input parameter **communicative goal**, which stipulates what is an acceptable level of belief, e.g., *BELIEVED* or *RATHER BELIEVED*.)

B. Otherwise, do nothing.

(The inference has had some effect on the proposition to be conveyed, but this effect is not sufficient to determine that the proposition is known by the addressee.)

ii. If an inference is correct, but does not correspond to a proposition in the list of propositions to be conveyed, then do nothing.

(The inference has no effect on the discourse[3].)

iii. If an inference is incorrect, then

If the belief in the affected proposition has fallen below the requirements established by the given communicative goal, then

Add the proposition to the list of propositions to be conveyed.

(Note that if the proposition in question was previously marked as deleted (Step i), it will be reinstated as a proposition that must be conveyed.)

(b) Update the student model with the above inferences.

(c) If the updated list of propositions to be conveyed is not empty, then

Add the RDs produced by Propose-RDs(*updated-list-of-propositions-to-be-conveyed, aspect*) to the RD proposed in this alternative.

To illustrate the workings of this algorithm, let us return to our Bracket Simplification example. For our discussion, we assume that the student is able to understand abstract explanations, i.e., L_{AU} and C_{AU}, the likelihood and confidence factor of the rule Abstract-Understand, are high. Now, the aspects to be conveyed with respect to Bracket Simplification are domain and operation. In the current implementation, we select operation first, since the inferences from the RDs generated to convey this aspect tend to affect other propositions to be conveyed. Next, we apply rules of inference in backward reasoning mode to generate RDs that can convey the proposition [Bracket-Simplification use-1 $+/-$]. This step yields the RDs {Assertion} and {Assertion + Instantiation}, where an {Assertion + Instantiation} contains an Assertion complemented with an Instantiation of the predicate in the asserted proposition. In our example, this predicate refers to the first step of Bracket Simplification. Both of these RDs have a high likelihood of conveying the intended proposition with a degree of belief that meets the input requirements given to the system. In both alternatives, the relationship use-1 in the Assertion is conveyed by a descriptor such

[3] Zukerman (1990a) describes a mechanism which produces discourse that addresses such inferences if they are weak.

as "before multiplying" which identifies the position of the $+/-$ operation in the Bracket Simplification procedure. The generation of this descriptor is performed by a procedure which generates referring expressions for the predicates in the propositions. Since the second step of Bracket Simplification is known and will not be described, the procedure proposes the descriptor "before multiplying." If both steps of Bracket Simplification are described, then ordinal conjunctive expressions, such as "first" and "second," are more suitable.

Let us now consider the alternative initiated by {Assertion}. In this case, the application of the inference rules in forward reasoning mode does not affect any of the other propositions to be conveyed. Hence, we update the student model to reflect the fact that the student has been informed of the first step of Bracket Simplification, and re-activate our algorithm with respect to the propositions in the aspect domain.

During the backward reasoning step, our mechanism determines that the RDs {Assertion} and {Assertion + Instantiation} may be used to convey the proposition [Bracket-Simplification apply-to Like-Terms]. In both cases, during the forward reasoning stage, the following inferences may be drawn from the Assertion: (1) a similarity-based inference based on the user's belief that Like Terms are similar to Numbers; (2) a generalization based on the belief that Like Terms are a subset of Algebraic Terms; and (3) a uniqueness implicature. The similarity-based inference corroborates the user's correct belief that Bracket Simplification applies to Numbers; the generalization corroborates his/her incorrect belief in the applicability of Bracket Simplification to Algebraic Terms; and the uniqueness implicature concludes that Bracket Simplification applies *only* to Like Terms, and hence not to Numbers or to Algebraic Terms.

The uniqueness implicature, which conflicts with the similarity-based inference and with the user's correct belief that Bracket Simplification applies to Numbers, is prevented by prefixing the proposed Assertion with a Mention that corroborates the user's belief, e.g., "*In addition to Numbers*, Bracket Simplification applies to Like Terms." At first glance, it appears that information that was omitted in the filtering process (Section 4.1.2) is now being reinstated. However, the generation of this preamble links the Assertion to an existing belief held by the user, rather than presenting this belief as if it were new information.

The generalization, which conflicts with the uniqueness implicature and corroborates the user's erroneous belief that Bracket Simplification applies to Algebraic Terms, is already being addressed by the second domain proposition in Table 5, which was proposed to contradict the erroneous belief. Hence, nothing needs to be added to the list of propositions to be conveyed. However, the fact that the generalization can be inferred from the proposed Assertion enables the system to record an expectation violation relation between the Assertion and the RD(s) that will be generated to convey the second domain proposition in Table 5. This relation, and other relations inferred in a similar way, are used to guide the discourse organization process (Zukerman and McConachy, 1993), and to generate appropriate Meta Comments which link the RDs in question (Zukerman, 1991). For instance, if RDs that contradict each other are juxtaposed in the discourse, Meta Comments such as "but" or "however" are suitable. However, if these RDs are separated by other RDs, more explicit Meta Comments, such as "Despite X," will be generated.

The generation of RDs that convey the second domain proposition in Table 5 is performed as described above, yielding {Negation} and {Negation + Instantiation}. This results in the output in Table 6 for the alternative where an Assertion was generated for the first and third proposition in Table 5, and a Negation for the second proposition. Our current implementation produces the rhetorical actions of the RDs and the propositional representation. The English text has been added for illustrative purposes.

Mention	[Bracket-Simplification apply-to Numbers] "In addition to Numbers,
Assert	[Bracket-Simplification apply-to Like-Terms] Bracket Simplification applies to Like Terms,
Negate	[Bracket-Simplification (always)apply-to Algebraic-Terms] but it does not always apply to Algebraic Terms.
Assert	[Bracket-Simplification use-1 +/−] In Bracket Simplification, we add or subtract the terms inside the brackets before multiplying."

Table 6: Sample Set of RDs Generated by the Content Planning Process

We conclude this discussion by describing briefly the alternative headed by {Assertion + Instantiation} of the proposition [Bracket-Simplification use-1 +/−]. The Instantiation of this proposition with respect to Like Terms, such as $3(2x + 5x)$, results in discourse which is markedly different from the discourse in Table 6. This is due to the fact that in the forward inference step, the generalization inference rule produces the inference [Bracket-Simplification apply-to Like-Terms] from this Instantiation. As a result, this proposition is deleted from the list of propositions to be conveyed. Table 7 depicts the output generated for this alternative.

Assert + Instantiate	[Bracket-Simplification use-1 +/−] $3(2x + 5x) = 3(7x)$ "In Bracket-Simplification, we add or subtract the terms inside the brackets before multiplying, e.g., $3(2x + 5x) = 3(7x)$.
Negate	[Bracket-Simplification (always)apply-to Algebraic-Terms] Notice that Bracket Simplification does not always apply to Algebraic Terms."

Table 7: Alternative Set of RDs Generated by the Content Planning Process

At this point in the content planning process, we have a number of candidate sets of RDs, where each set conveys the specified aspects of the intended concept. If a particular set of RDs contains a concept that is not well understood by the hearer, WISHFUL generates subordinate sets of RDs, each of which is an alternative way of conveying this concept (Section 4.3). Upon completion of this step, the alternative with the least number of RDs is selected, which in our example is the alternative in Table 7.

4.3 Conveying the Concepts in each RD

In order to ensure that the user understands all the RDs in a set of RDs, the content planner performs the following actions for each of the concepts mentioned in a set of RDs: (1) it determines the aspects of the concept which are relevant to the understanding of the propositions which contain the concept, (2) it determines a communicative goal for these aspects, and (3) it regresses to generate RDs that accomplish this communicative goal with respect to the selected aspects of the concept.

The determination of the aspects the user must know about a concept in order to understand a proposition which contains this concept is based on the main predicate of the proposition and on 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] proposed above, the user must know what Like-Terms are and what they look like. Hence, the system returns the aspects *membership-class* and *structure*.

The determination of a communicative goal with respect to the selected aspects of a concept is based on the relevance of this concept to the concept given originally as input to the system. That is, the more relevant a concept in an RD is to the original concept, the better it should be known by the user. This consideration is implemented by lowering the expertise requirements with respect to a concept as the recursion becomes deeper. In this manner, we preclude the elaboration of concepts which are far removed from the main concept to be conveyed, while at the same time, ensuring a minimal level of competence with respect to these concepts.

5 Future Research

Several aspects of the student model stand out as candidates for future work. As indicated in Section 3.1, the initial level of detail of the material is determined by the designers of the knowledge base. An interesting enhancement to our knowledge representation scheme involves the implementation of a mechanism that automatically collapses a node representing complex information and its related satellites into a single node when the student has a sufficient grasp of the concepts involved. This is equivalent to a student learning to perform a complex procedure, e.g., Bracket Simplification or Short Division, in a single step, rather than in a sequence of discrete component steps, as s/he would if s/he was unfamiliar with the procedure. Conversely, if the initial level of detail is too difficult for a particular student, a mechanism that expands already collapsed nodes is also required.

A second area of interest involves extending and refining the set of inference rules supported by the system. For example, the following inference rule may be incorporated into the set of inference rules:

{Known: [A isa B] AND [C isa D]} AND {**Assert** [B isa C]} ⇒

Either D is an ancestor of A or there is an inconsistency.

This rule checks the consistency of the knowledge which results from the combination of newly asserted information with existing beliefs. As stated in Section 3.2.2, another possible refinement of our inference rules consists of incorporating typicality considerations into them to reflect the fact that unsound inference rules which are applied to items that are typical of a group are more likely to yield correct results than unsound rules which are applied to atypical items.

At present, our mechanism generates stand-alone explanations. However, when this mechanism is incorporated into a fully interactive system, the different aspects of our student model must be verified and possibly adjusted after interacting with a student. An interesting avenue of investigation involves examining the effect of the system's output on the student's confidence in his/her acquired beliefs. For example, it is possible that after an extended interaction with WISHFUL, the student will trust the system to dispel all his/her erroneous inferences. In this case, any inference that is not dispelled (including erroneous ones) may be strongly believed by the student. Another interesting line of investigation for such a system consists of activating the inference rules in a reflective mode after a session with a student has been completed. In this mode, the system would draw further reaching inferences from the generated discourse. Typically, these inferences would interact with each other, thereby requiring a processing mechanism that combines the inferences until the beliefs in the student model reach quiescence. The result of this process would then be the starting point of the next interaction with the student.

Finally, Propose-RDs is a cautious procedure rather than an optimal one. That is, it presents at least as much information as is necessary to achieve the given communicative goal. However, owing to its sequential operation, it may miss opportunities to omit superfluous information. This happens when an RD generated at a later stage conveys indirectly information for which an RD was generated earlier. This situation is addressed in part by our heuristic for selecting the aspects to be conveyed, but this heuristic does not guarantee optimality. A mechanism which optimizes the output of the content planner is currently under investigation.

6 Conclusion

The content planning mechanism presented in this paper generates RDs by consulting a student model which represents three aspects of a student, namely his/her beliefs and skills, the inferences s/he is likely to draw from the presented information, and his/her ability and attitude. Our mechanism improves the conciseness of the generated discourse by omitting information that is known by the student or which the student can easily infer from the discourse. In addition, our mechanism addresses a student's erroneous beliefs and his/her possible incorrect inferences from the discourse.

Our mechanism has been implemented in Sun Common Lisp on a SPARCstation 2. It takes less than 1 second of CPU time to execute the example discussed in the paper and other examples in high school algebra involving several RDs.

The student model consulted by our content planning mechanism must represent accurately the beliefs and skills of individual students. Such a model may be acquired with the help of a diagnostic system, such as those described in (Sleeman, 1982; Burton, 1982). In contrast, uniqueness implicatures are influenced by expectations which are common to all users, in addition to the wording of the discourse. The inference rules are also commonly applied by all types of students, but the conditions for the application of the direct and indirect inference rules and for the acquisition of the conclusions they draw vary for different types of students.

Since at present our system is not interactive, a preliminary validation of our student model and our content planning mechanism was performed as follows: WISHFUL was activated with the different student profiles, and its output was translated manually into English. We then showed these translations to tutors and lecturers in the Department of Computer Science at Monash University, and to students, both at university and in school. The tutors and lecturers were asked to select texts that best suited certain types of students, e.g., they had to select which text they would show to an *EXCELLENT* student, a *MEDIOCRE* student, etc, and they also had the option to indicate that a text was not suitable for any student. The students were shown all the texts, and were asked to select the text they thought was the clearest. There was general agreement between the discourse planned by WISHFUL for particular types of students and the texts selected by the interviewed students and teaching staff. Further, comparison of WISHFUL's output with texts found in introductory textbooks showed that WISHFUL's output for *MEDIOCRE* students is similar in content to the material found in these textbooks. Once the system generates textual output, a more complete evaluation will be performed by showing the generated texts to significant populations from the various target audiences, and comparing their response to the text generated by WISHFUL with their response to texts from algebra textbooks.

References

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

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

Buchanan, B., and Shortliffe, E.: 1985, *Rule-based Expert Systems: The MYCIN Experiments of the*

Stanford Heuristic Programming Project, Addison-Wesley.

Burton, R.R.: 1982, 'Diagnosing Bugs in a Simple Procedural Skill'. In *Intelligent Tutoring Systems*, Sleeman, D. and Brown, J.S. (Eds.), London: Academic Press, pp. 157-183.

Carberry, S.: 1988, 'Modeling the User's Plans and Goals'. In *Computational Linguistics* **14**(3), pp. 23-37.

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

Cawsey, A.: 1991, 'Using Plausible Inference Rules in Description Planning'. In *Proceedings of the Fifth Conference of the European Chapter of the ACL.*

Clark, D. A.: 1990, 'Numerical and symbolic approaches to uncertainty management in AI'. In *Artificial Intelligence Review* **4**, pp. 109-146.

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

Driankov, D.: 1986, 'Towards a Many-Valued Logic of Quantified Belief'. Doctoral Dissertation, Department of Computer and Information Science, Linköping University, Sweden.

Dyer, M.G.: 1983, *In-Depth Computer Understanding: A Computer Model of Integrated Processing for Narrative Comprehension*, MIT Press, Cambridge, Massachusetts.

Green, N., and Carberry, S.: 1992, 'Conversational Implicatures in Indirect Replies'. In *Proceedings of the Thirtieth Annual Meeting of the Association for Computational Linguistics*, Newark, Delaware, pp. 64-71.

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

Grice, H.P.: 1978, 'Further Notes on Logic and Conversation'. In *Syntax and Semantics, Volume 9: Pragmatics*, Cole, P.J. and Morgan, J.L. (Eds.), Academic Press.

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

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

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

Joshi, A., Webber, B.L., and Weischedel, R.M.: 1984, 'Living Up to Expectations: Computing Expert Responses'. In *AAAI-84, Proceedings of the National Conference on Artificial Intelligence*, Austin, Texas, pp. 169-175.

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

Litman, D. and Allen, J.F.: 1987, 'A Plan Recognition Model for Subdialogues in Conversation'. In *Cognitive Science* **11**, pp. 163-200.

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

Maybury, M.T.: 1990, 'Custom Explanations: Exploiting User Models to Plan Multisentential Text'. In *Proceedings of the Second International Workshop on User Modeling*, Honolulu, Hawaii.

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

Meteer, M.: 1991, 'Bridging the Generation Gap between Text Planning and Linguistic Realization'. In *Computational Intelligence* **7**(4), Special Issue on Natural Language Generation, pp. 296-304.

Moore, J.D.: 1989, 'A Reactive Approach to Explanation in Expert and Advice-Giving Systems'. Doctoral Dissertation, Department of Computer Science, University of California, Los Angeles, California.

Moore, J.D. and Paris, C.L.: 1992, 'Feedback to Compensate for the Unreliability of User Models'. In *User Modeling and User-Adapted Interaction* **2 (4)**, pp. 287-330.

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

Norvig, P.: 1989, 'Marker Passing as a Weak Method for Text Inferencing'. In *Cognitive Science* **13**, pp. 569-620.

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

Reiter, E.: 1990, 'Generating Descriptions that Exploit a User's Domain Knowledge'. In *Current Research in Natural Language Generation*, Dale, R., Mellish, C. and Zock, M. (Eds.), Academic Press, pp. 257-285.

Sleeman, D.: 1982, 'Inferring (mal) Rules from Pupil's Protocols'. In *ECAI-82, Proceedings of the European Artificial Intelligence Conference*, pp. 160-164.

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

Suthers, D.: 1991, 'A Task-appropriate Hybrid Architecture for Explanation'. In *Computational Intelligence* **7**(4), Special Issue on Natural Language Generation, pp. 315-333.

van Beek, P.: 1987, 'A Model for Generating Better Explanations'. In *Proceedings of the Twenty-Fifth Annual Meeting of the Association for Computational Linguistics*, Stanford, California, pp. 215-220.

van Lehn, K.: 1983, 'Human Procedural Skill Acquisition: Theory, Model and Psychological Validation'. In *AAAI-83, Proceedings of the National Conference on Artificial Intelligence*, Washington, D.C., pp. 420-423.

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

Zukerman, I.: 1990b, 'Consulting a Shallow Model of a Listener to Generate Supportive Rhetorical Devices'. In *Proceedings of the Second International Workshop on User Modeling*, Honolulu, Hawaii.

Zukerman, I.: 1991, 'Using Meta Comments to Generate Fluent Text in a Technical Domain'. In *Computational Intelligence* **7**(4), Special Issue on Natural Language Generation, pp. 276-295.

Zukerman, I., and McConachy, R.: 1993, 'An Optimizing Method for Structuring Inferentially Linked Discourse'. In *AAAI-93, Proceedings of the National Conference on Artificial Intelligence*, Washington, D.C.

Contents

List of Figures

List of Tables