LISTENER MODEL FOR THE GENERATION OF META-TECHNICAL UTTERANCES IN MATH TUTORING

Ingrid Zukerman
Judea Pearl

December 1984
Report No. CSD-840064

Listener Model for the Generation of Meta-Technical Utterances in Math Tutoring

Ingrid Zukerman and Judea Pearl

Computer Science Department
University of California, Los Angeles
Los Angeles, CA 90024
telephone number: (213) 825-3243
netmail address: ingrid@UCLA-LOCUS.ARPA
judea@UCLA-LOCUS.ARPA

Topic: Natural Language

Keywords: language generation, tutoring systems, meta-technical utterances, cognitive model

Word count: 4281

Abstract

When reading or listening to an explanation of a technical subject, we notice the presence of conversational expressions like "however," "as I have stated before," "next," "generally speaking," etc. These expressions are not included in the discourse merely for decorative purposes, but carry important information which the listener uses to speed up the comprehension process. In this paper we model the meaning of these expressions in a representation that facilitates the generation of fluent and cogent discourse. As a testbed for our ideas we have selected the domain of high-school algebra; specifically, we are implementing a system called FIGMENT, which generates commentaries on the solution of algebraic equations.

1. Introduction

When generating tutorial text in a technical subject, a teacher wishes to present the information in a manner which is most accessible to the student. Clearly, a necessary precondition is that the teacher transmit the appropriate information items. However, we also notice the presence of expressions like

This work was supported in part by the National Science Foundation Grant IST 81-19045

"however," "as I have stated before," "next," "generally speaking," etc., which are *not* part of the subject matter. These expressions, denoted Meta-Technical Utterances (MTUs), carry important information which assists the listener in assimilating the knowledge being transferred.

Previous research on the semantics of a subset of these utterances (Farnes 1973, Winter 1968 and 1977, Reichman 1978 and 1984 and Hoey 1979) indicates that the presence of an MTU can *signpost* what kind of information is to be presented in the forthcoming sentence or sentences. Furthermore, Farnes claims that "the identification and use by readers of such cues, greatly aids comprehension," and Hoey points to the fact that problems of comprehension have been shown to arise due to "faulty" or missing signaling.

The text generated by natural language generation systems designed by Davey (1979), Mann and Moore (1980), McKeown (1982), Swartout (1982) and Kukich (1984) contain mostly MTUs like "however," "next," "therefore," "or," etc., which directly reproduce the speaker's organization of the subject matter, i.e., they represent the relationship between two or more items of knowledge. For example, if item B violates the expectations established by item A, this relationship is expressed by the utterance: "A, however B." This type of MTUs shall be denoted *Knowledge-Organization MTUs*.

These, however, account only for a fraction of the MTU's found in natural discourse. Teachers of technical subjects often use MTUs like "as I have stated before," "let us try another approach," "in other words," "this method is somewhat complicated," etc. These utterances are more intimately connected with the listener's learning process than with the organization of the subject matter. In the research reported in this paper, we attempt to gain insight into this learning process, by building and implementing a generative model of the meaning of these MTUs.

In the following section we shall present a functional classification of both types of MTUs. Then we shall describe the system that generates them.

2. Functional Taxonomy of MTUs

The classification of Meta-Technical Utterances presented in this section is based on their function, as seen by the tutor, in transmitting the subject matter to the student. In our taxonomy we recognize three main functions of MTUs: (1) Knowledge Organization, (2) Knowledge Acquisition and (3) Affect Modification. In the following subsections we shall consider in detail the MTUs that perform these functions.

2.1 Knowledge-Organization MTUs

At each point in time, the information residing in the student's (or tutor's) mind can be visualized as a network whose nodes contain the individual information items, and whose links contain the relations between the nodes. For example: NODE1 contains the *purpose* of NODE2, or NODE3 is an *alternative* to NODE1. These relations directly reflect the knowledge about equations and their solution (as opposed to the state of the dialogue or the relationship between the teacher and the student), and they roughly correspond to Hallyday and Hasan's *external* category (in Hallyday and Hasan 1976) and Longacre's *basic* heading (in Longacre 1976). The Knowledge-Organization function consists of transmitting these relations, and is performed by the following types of MTUs:

Additive - these MTUs signal *additional events* [and, also, nor]; *realization of expectations* [indeed] or *availability of additional alternatives* [alternatively, or, or else],

Adversative - these MTUs signal *violation of expectations* [however, nevertheless, although, but, despite this]; *dismissal,* which is triggered when two different paths in a solution arrive at the same pattern [in any case, either way, at any rate] or *avowal,* which is triggered when a well known pattern, implicit in the equation, is recognized and made explicit [notice that, actually, as a matter of fact, in fact],

Causal - these MTUs signal the *reason for performing an operation* [for this reason, on account of this, it follows, therefore, so, then, because of this]; its *purpose* [for this purpose, with this in mind, to this end]; *hopes* [hopefully, expecting to get]; *result* [as a result, in consequence]; *means* [this can be accomplished by], or *correctness* [this works because], and

Temporal - these MTUs signal *sequence* [then, next, after that, first ... next, finally, previously] or *partial*

sequence [at the same time].

Whereas the values of these relations would in general depend on the solution found for each particular equation, most of the relations in the Causal category do not vary from equation to equation, but are inherent to the solution methods. For example, each method has a purpose, a pattern or reason for applying it, a correctness proof, etc. Thus, the knowledge required for commenting on the causal aspects of the solution of equations should preferably reside in a module that contains the expertise about strategies for solving equations. We call this module a Problem-Solving Expert (see section 3).

2.2 Knowledge-Acquisition MTUs

The MTUs that perform the Knowledge-Acquisition function are related to the interaction between the teacher and the student, rather than to the subject matter itself. They ease the assimilation of the subject matter by providing information *about* it. However, once the material has been mastered, most of this information is not remembered by the student. In the context of tutoring algebra, the Knowledge-Acquisition function is performed by the following types of MTUs:

Motivational - a teacher will often use this type of MTUs to motivate the student to listen to the forthcoming technical utterance. For example, if a new method is to be taught, the tutor might say: "This method is very expeditious," and if the student has to practice the same type of equation many times, the teacher might say: "Third degree equations are rather difficult and demand lots of practice."

Focal - a student generally attempts to process a forthcoming technical utterance in the currently active focus space (Grosz 1977). If the teacher wants the student to change the active focus space, or to close an open focus space, he needs to present the student with an MTU to this effect. For example, the Focal-MTU "Let us consider the following equation:" will close the focus space corresponding to the previous equation, and open and activate a new focus space for the next equation. Temporary focus shifts are signaled by MTUs like "incidentally" or "by the way."

Categorical - this type of MTUs is used by a teacher to clarify the relationship between the forthcoming technical utterance and a specific item in active focus. For example, a tutor might say: "Let's take

the first term on the right hand side, **namely** $x^2(x - 3)$, ..." In this example, the MTU "namely" alerts the listener to the fact that the explicit term merely paraphrases the preceding positional description and is not to be taken in addition to the first term on the right hand side. Categorical MTUs are included in the taxonomies presented in Hallyday and Hasan's *internal* classification (in Hallyday and Hasan 1976) and in Longacre's *elaborative* heading (in Longacre 1976). MTUs which belong to this subclass are: "in other words," "generally speaking," "for example," etc.

Implementational - these MTUs prepare the student to select a computational activity required for assimilating the technical utterance that follows. We have identified two main types of activities: *adding* an item to one's knowledge pool, and *verifying* the workings of existing knowledge (for possible revision). For example, if the tutor wishes the student to use the forthcoming technical utterance to verify existing knowledge, he should signal his intent by means of an MTU like "as I have stated before." On the other hand, if the teacher wants the student to prepare for learning a new subject (i.e., transfer to addition mode) he might say: "We shall now introduce a new type of equation."

Estimational - these MTUs inform the student that the forthcoming technical utterance is of unusual length and/or complexity. Examples are: "This equation is rather straightforward," or "The following method is somewhat lengthy."

In order to illustrate the importance of these Knowledge-Acquisition MTUs, let us examine the following imperfect discourse:

"*The method of completion to square for solving quadratic equations works as follows*"

... description of completion to square method ...

... examples of application of this method ...

"*In order to perform completion to square we have to execute the following steps*"

... description of completion to square method ...

The dissonance in the preceding discourse stems from the repetition of preparatory directives in lines 1 and 4. The first time the completion to square method is introduced, the MTU "*The completion to square ... works as follows*" performs the dual role of establishing our expectations for receiving new

knowledge and of signaling that the forthcoming technical utterance is a description of a method already in active focus (see Causal subclass in Knowledge-Organization, above). The second time (fourth line), the MTU repeats the directives given in the previous MTU (even though their English representations differ). We again prepare ourselves to add the method description to our knowledge pool. But upon discovering that this information is already in our knowledge pool, we experience the discomforting feeling that the discourse generator has no recollection of previous activities. Notice, however, that had the fourth line been replaced by: "*Let's go over the method of completion to square again*," the learning process would be much smoother; we would prepare to use the incoming information for verifying already existing knowledge, instead of treating it as new information to be added on.

This example illustrates the claims made by Hoey, Farnes and others (see Introduction, above), regarding the adverse effects of improper usage of MTUs on the learning process. In order to generate a commentary which includes Knowledge-Acquisition MTUs, we need a module that represents how both technical and meta-technical utterances influence the listener's mental activities. This module would then inspect the technical utterances about to be issued, determine their effect on the listener, and generate adequate Knowledge-Acquisition MTUs. We shall call this module a Cognitive-Processes Module (see section 3).

2.3 Affect-Modification MTUs

One of the tutor's goals is to teach the student which algebraic operations and results are considered favorable and which are considered unfavorable. In addition, the tutor wishes the attitude of the student to remain positive throughout the session. To achieve these goals, a tutor may be required to use Affect-Modification MTUs, which we divide into two subclasses according to their goals:

Affect-transference - if the tutor is of the opinion that the forthcoming technical utterance should have an affective impact on the student, he might precede it by an MTU like "Unfortunately" or "Fortunately" (according to the situation). For example, if the teacher says: "Unfortunately, the only way of solving this equation is to open parentheses and collect terms," the student should understand that this method of solution is considered undesirable and should only be used as a last resort. This point might be missed by the student, had the MTU "unfortunately" been omitted.

Consolation - one of the teacher's goals is to maintain a positive student attitude throughout a tutorial session. This goal can be partially accomplished by using Knowledge-Acquisition MTUs, as exemplified in the scenario presented in section 2.2. There exist, however, situations in which negative affects cannot be prevented by means of these MTUs. For example, a student may fail to understand a solution method, in spite of having received preparatory Knowledge-Acquisition MTUs. In cases like this, a teacher should reassure and console the student, so that he will be able to continue learning. This is the purpose of Consolation MTUs such as: "This was a very difficult exercise. Let's now try something else" or "Don't worry, I will explain this a few more times."

Consolation MTUs are also required when a technical utterance carries upsetting information. For instance, if a particular method fails to solve an equation, the tutor may say: "Fortunately, however, there is another way of solving this equation."

The reader will notice that some of the Affect-Modification MTUs are equivalent to or contain Knowledge-Acquisition MTUs (in particular Estimational-MTUs). The reason for this is that an MTU may perform several illocutionary acts (Appelt 1982), e.g., it can signal to the student which mental resources to prepare and, simultaneously, dispel negative affects.

To generate commentaries which have a desired affective influence on a student, a discourse generator needs a module which recognizes the affective impact of both technical and meta-technical utterances. This module will then neutralize anticipated negative affects by generating adequate Affect-Modification MTUs.

3. Design of FIGMENT

The system outlined in this section is designed to generate a fluent and cogent commentary on algebraic equations, based on the taxonomy presented in the preceding section. The generation of the commentary is performed in three stages:

i. In the first stage, a **technical message** is decided upon. This task is performed by means of three modules: (1) Problem-Solving Expert, (2) Student Model and (3) Tutoring Strategist. Since these modules are usually part of intelligent tutoring systems (see articles by Clancey, Genesereth, Gold-

stein, and others in Sleeman and Brown 1982), we shall discuss them only briefly. The Problem-Solving Expert solves the equation and outputs a graph in which each branch contains an attempted solution alternative. Next, the Tutoring Strategist modifies this graph by suppressing alternatives and steps which are well known to the student and adding explanations where necessary (e.g., purpose of an operation, its description, etc.). In order to perform their task, both these modules use information about the state of the student's knowledge; this information resides in the Student Model. The modified graph produced by the Tutoring Strategist (see fig. 3.1) is used directly for generating Knowledge-Organization MTUs, since it contains all the information necessary for identifying the relationships mentioned in section 2.1.

ii. In the next stage, the discourse is complemented and revised by adding Knowledge-Acquisition and Affect-Modification MTUs. The Knowledge-Acquisition MTUs are added by the Cognitive-Processes Module, which simulates some of the problem-solving strategies and comprehension processes used by the student. The Affect-Modification MTUs are generated by a unit called an Affects Module, which models the tutor's handling of affects connected with the learning situation.

iii. In the final stage, a Sentence Composer organizes the completed message into paragraphs and sentences, and a Grammar translates them into English.

The following subsections describe the Cognitive-Processes Module and the Affects Module.

3.1 Cognitive-Processes Module

The Cognitive-Processes Module purports to generate Knowledge-Acquisition MTUs for a technical message produced by the Tutoring Strategist. This message contains both technical information and an implicit representation of Knowledge-Organization MTUs.

As stated before, the role of Knowledge-Acquisition MTUs is to ease the assimilation of the subject matter by providing information *about* it. This information helps the listener prepare himself for processing the forthcoming technical utterance.

EQUATION: $(x-3)^2 - 4(x-3) - 12 = 0$

(ALTERNATIVE 1)
PATTERN: { x-3 common factor }
PRECONDITION: substitute x=3 in equation YIELD -12=0
[root of factor] ≠ [root of equation]
PRECONDITION NOT FULFILLED
STOP

(ALTERNATIVE 2)
ACTION: open parentheses, collect terms
RESULT: $x^2 - 10x + 9 = 0$
CONTINUE

(ALTERNATIVE 3)
PATTERN: { x appears only in expression x-3 }
CAUSE ACTION
1. ACTION: substitute y=x-3
PURPOSE: get canonic expression
RESULT: $y^2 - 4y - 12 = 0$

2. ACTION: factorization
RESULT: $(y-6)(y+2) = 0$

3. ACTION: solve for each factor
RESULT: y=6 or y=-2

4. ACTION: substitute back x-3 for y
RESULT: x-3 = 6 or x-3 = -2

5. ACTION: solve for each factor
RESULT: x=9 or x=1
SOLVED

Fig. 3.1: Output of the Tutoring Strategist

In order to generate these MTUs, the Cognitive-Processes Module simulates several processes which the listener undergoes upon hearing an utterance. The simulation proceeds in a manner similar to the analysis we performed on our example in section 2.2. If any of these processes results in negative affects, then a Knowledge-Acquisition MTU is prefixed to the utterance. The processes simulated correspond to the different types of Knowledge-Acquisition MTUs presented in section 2.2, namely: motivational, focal, categorical, implementational and estimational. Figure 3.2, for example, illustrates the process of establishing the implementation mode for the forthcoming technical utterance.

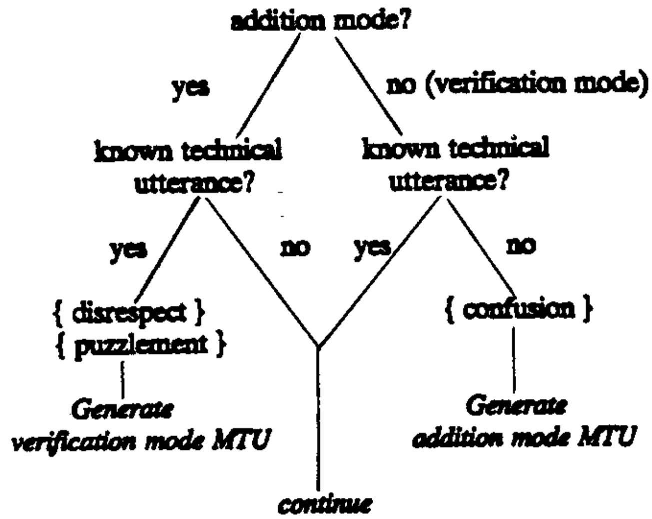

Fig. 3.2: Process for establishing implementation mode

According to this model, if a student is in the addition mode and receives a known technical utterance, he will experience puzzlement and disrespect (like in our scenario in section 2.2). However, if the tutor puts him first in verification mode (by means of an MTU like "As I have said before" or "Let's go over this one more time"), no adverse affects will result, and the learning process can continue. Alternatively, if the student is in verification mode and is presented with a new technical utterance, he will feel rather confused. In this case the teacher has to transfer him first to addition mode by using an MTU like "Let's now consider a new type of equation."

For each of the different types of the Knowledge-Acquisition MTUs, there exists a discrimination net similar to the one depicted above. A technical utterance traverses each of these nets in order to ascertain which MTUs it requires. At first, the input to these nets consists of the entire technical message, and FIGMENT determines whether it has to generate introductory MTUs related to the equation and its solution. A typical MTU of this kind would be the Focal-MTU: "Let us consider the following equation." Next, each of the major solution alternatives traverses these nets, causing the Cognitive-Processes Module to generate MTUs like (Motivational) "The following alternative is more elegant than the previous ones." Finally, the individual steps in each solution alternative traverse the discrimination nets, and, after generating the necessary MTUs, the Cognitive-Processes Module outputs an extended message which is fed to the Affects Module.

This process enables FIGMENT to generate MTUs that perform the same cognitive function (e.g., specifying implementation mode), but whose English representation differs depending on the type of the technical utterance under consideration. For instance, the MTUs "As I have stated before" and "This equation is similar to ..." can both be used to put a student in verification mode; however, while the former refers to explanations, the latter applies to equations.

3.2 Affects Module

The Affects Module completes the extended message by adding to it Affect-Modification MTUs as necessary.

Naturally, affects are classified as being either positive (satisfaction, pride, etc.) or negative (apprehension, confusion, etc.). There are two basic approaches for disposing of negative affects: they can either be avoided (by means of Knowledge-Acquisition MTUs), or, if they have already occurred, they should be dispelled and neutralized by positive affects. Sometimes this task can be accomplished by Consolation MTUs.

The Affects Module recognizes two types of situations in which a student might experience negative affects:

1. Failure to process a technical utterance - it is our contention that a student will fail to process a technical utterance if its relative difficulty and/or length is too high. Therefore, the Affects Module predicts a situation of failure by assessing the knowledge status and ability of the student against the difficulty and length of a technical utterance. If the result of this evaluation exceeds a certain threshold, a consolation MTU like "Don't worry, I'll go over this explanation a few more times" or "I know this equation was rather difficult" is generated.

2. Disturbing information carried by a technical utterance - there are cases in which the technical utterance itself carries information which may frustrate the student, such as when a third order equation is introduced for which no general solution shall be taught. At first, FIGMENT will test whether the forthcoming technical utterance neutralizes the effects of the disturbing one. If the result of the test is positive, the system will generate a Consolation MTU like "Fortunately, how-

ever, there is a way of solving this particular equation." If the test fails, then FIGMENT will generate an MTU of a more social nature, for example: "There is nothing we can do about this exercise, so let's try another problem."

An additional task performed by the Affects Module is the generation of Affect-Transference MTUs, which communicate the opinion of the tutor regarding the merit or lack of merit of an event, algebraic operation, etc. This task is accomplished by classifying technical information according to their affective impact into positive, negative and neutral, and then generating appropriate Affect-Transference MTUs like "Fortunately," "Unfortunately," etc.

After the relevant Affect-Modification MTUs have been generated, the message contains a list of technical utterances interleaved with code specifying requirements for MTUs. One of the tasks of the Sentence Composer is to collect this code into a stylistically sound representation which satisfies these requirements.

4. Conclusions

The distinctive feature of the design outlined in this paper is the generation of Meta-Technical Utterances by consulting simplified models of the cognitive processes activated by the listener. We have shown that, at least in principle, a text generation process based on such models would capture most rhetorical features that support natural discourse. It is generally believed that any system which generates continuous discourse must contain models representing both the process by which a listener absorbs information and the impact of this information on the listener's affects. This paper offers a concrete design of the makeup of these models and their usage in text generation.

References

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

Davey, A. (1979), Discourse Production. Edinburgh University Press, Edinburgh.

Farnes N.C. (1973 revd. 1975), Comprehension and the use of Context, Unit 4, Reading Development. Educational Studies: a Post-Experience Course and 2nd Level Course P.E. 261.

Grosz B.J. (1977), The representation and use of focus in dialogue understanding. Ph.D. dissertation, University of California, Berkeley.

Hallyday M.A.K. and Hasan R. (1976), Cohesion in English. Layman Press, London.

Hoey M. (1979), Signaling in Discourse. English Language Research, University of Birmingham, Birmingham Instant Print Limited.

Kukich K. (1984), The Feasibility of Automatic Natural Language Report Generation, Ph.D. Dissertation, Carnegie-Mellon University, Pittsburgh.

Longacre, R. E. (1976), An Anatomy of Speech Notions. Peter de Ridder Press Publications in Tagmemics No. 3.

Mann, W.C. and Moore, J.A. (1980), Computer as Author - Results and Prospects. Report No. ISI/RR-79-82, Information Sciences Institute, Los Angeles, January 1980.

McKeown, K.R. (1982), Generating Natural Language Text in Response to Questions About Database Structure. Doctoral Dissertation, The Moore School of Electrical Engineering, University of Pennsylvania, Philadelphia.

Reichman R. (1978), Conversational Coherency. Cognitive Science 2, pp. 283-327.

Reichman-Adar R. (1984), Extended person-machine interface. Artificial Intelligence 22, pp. 157-218.

Sleeman D. and Brown J.S. (Eds.) (1982), Intelligent Tutoring Systems, London: Academic Press.

Swartout W.R. (1982), XPLAIN: A System for Creating and Explaining Expert Consulting Programs, USC/Information Sciences Institute.

Winter E. O. (1968), Some Aspects of Cohesion. In Sentence and Clause in Scientific English, by R. D. Huddleston et al., Communication Research Centre, Department of General Linguistics, University College, London, May 1968, pp. 560-604.

Winter E. O. (1977), A clause-relational approach to English texts: A study of some predictive lexical items in written discourse. Instructional Science, Vol. 6, No. 1, January 1977, pp. 1-92.