

Expectation Verification: A Mechanism for the Generation of Meta Comments

Ingrid Zukerman
Department of Computer Science
Monash University

ABSTRACT

Meta Comments, such as “however,” “as I have stated before” and “if you pardon the expression,” are pervasive in human discourse. In this paper, we present a predictive mechanism for the generation of Meta Comments based on the tenet that they signal a change in the status of beliefs and expectations a listener is conjectured to possess. In particular, our mechanism anticipates a listener’s expectations by activating prescribed inferences on the state of the discourse. It is being implemented in a system called FIGMENT II which generates commentaries on the solution of algebraic equations.

INTRODUCTION

It is generally believed that an important aspect of competent writing involves anticipating and addressing a reader’s[†] beliefs and expectations arising from the presented information. This premise has been applied to a certain extent in natural language parsing (Riesbeck 1982), however, it has been largely ignored in natural language generation systems. Text generation systems implemented to date focus on the generation of utterances which directly reflect a speaker’s communicative intent, under the implicit assumption that the only inferences drawn by a listener are those which directly follow from the presented information (Appelt 1982, Swartout 1982, Kukich 1984, McKeown 1985, Paris 1987, Hovy 1987).

In this paper, we consider the generation of a class of Meta Comments (MCs), denoted *Expectation Verification MCs*, which are generated to notify a reader of an impending change in the status of his/her beliefs or expectations due to a forthcoming message. We postulate that these MCs are generated, if, in their absence, the reader is likely to experience affective responses which may inhibit the acquisition of forthcoming information. In particular, we focus on the generation of MCs with respect to the attributes of a message, such as focus, length and style. For instance, if a normally strict math teacher walks into the classroom and starts talking about a movie he saw recently, most students are likely to wonder about this abnormal behaviour, instead of attending to the current discourse. This adverse effect may have been minimized or avoided by means of an MC indicative of a shift from the expected focus, such as “Before we start with equations, I would like to tell you about” (If this behaviour constitutes an extreme departure from the norm, additional rhetorical devices, such as justification or motivation, may be called for.)

In the next section, we consider beliefs and expectations which affect the generation of MCs. We then examine a mechanism for generating MCs and discuss the different types of MCs accounted for by this mechanism.

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

EXPECTATIONS AND BELIEFS

During the knowledge acquisition process, a listener exhibits expectations with respect to various aspects of a message, such as its content, level of difficulty, length and context. A correctly understood message may still inhibit the smooth continuation of the knowledge acquisition process, if either (1) the listener fails to notice its effect on the status of his/her expectations or beliefs, or (2) its effect is noticed, but due to the absence or inadequate use of an MC, the listener loses confidence either in the speaker's ability, as in the classroom scenario discussed above, or in his/her own perception.

The beliefs and expectations which are relevant to the text generation task in general, and the generation of Meta Comments in particular, are those which a speaker conjectures are maintained by a listener. We distinguish between two types of beliefs/expectations based on the type of information they refer to:

- i. *Content Beliefs* pertain to a listener's knowledge. In particular, *Content Expectations* usually pertain to the result of an action or a feature of an object, and are largely domain dependent. In a technical domain, common content expectations are for: (1) the existence/absence of a solution, (2) a particular result, and (3) a particular sequence of steps in a solution path.
- ii. *Attribute Beliefs* pertain to meta-knowledge, i.e., attributes of an item of knowledge which affect its acquisition, such as (1) Focus, (2) Informational Status (new or previously seen?), (3) Length, (4) Difficulty, (5) Style, (6) Timing, and (7) Orientation with respect to a baseline measure. In general, attribute beliefs are relevant to discourse generation when they become *Attribute Expectations* with respect to a forthcoming message.

The status of an expectation may be *active*, *inactive* or *confirmed*. An active expectation may be inactivated by its violation or realization. However, when an active expectation for some result requires the existence of a sequence of events or objects which are necessary to produce this result, each of the elements in the sequence has to agree with the expectation. In this case, if the first element in the sequence complies with the expectation, its status is changed to confirmed. This status is maintained until the expectation is either violated or fully realized. The status of a belief may change due to the presentation of contradicting or corroborating information (violation and realization, respectively).

A change in the status of an expectation or a belief may be accomplished *explicitly*, by means of a direct inference from an utterance, or *implicitly*, by means of an indirect inference. In addition, we postulate that people also use some general purpose inference mechanisms to activate implicit attribute expectations. We model these mechanisms by means of *Expectation Activation Demons (EADs)*. The following Expectation Activation Demons account for a variety of attribute expectations maintained by a listener.

Inertia reflects one's tendency to expect the continuation of an established behaviour pattern.

Default represents expectations associated with a context. This context may be established by the discourse focus, the social setting and/or the speaker.

Diversity represents an expectation for a forthcoming utterance to be different from a recently presented or known utterance. The level of discourse for which this expectation is active depends on the current context.

The activation of EADs may yield conflicting expectations. However, when more than one demon is applicable to a given situation, the expectation activated by the demon which yields the most specific expectation is accepted. In addition, the activation of conflicting expectations at the same level of specificity indicates that an MC should be generated to resolve the conflict.

GENERATING META COMMENTS

In order to determine the need for MCs, a system must keep track the evolution of a listener's conjectured expectations and beliefs. To this effect, for each message, it must activate direct and indirect inferences which affect a listener's attribute and content beliefs and expectations, and also activate EADs with respect to attribute expectations. Upon discovering that a message causes a change in the status of a belief or an active expectation, one or more MCs which indicate the nature of the change may be proposed. This process supports the generation of the following types of MCs.

- i. Expectative — Indicate a violation or realization (confirmation) of an expectation or a belief entertained by the speaker and/or the listener. They range from the implicit [*“However,” “Indeed”*] to the explicit [*“We were expecting X, but we got Y,” “As you may have suspected, the butler did it”*]. Expectation Violation MCs are also used to signal a violation of orientation expectations, e.g., *“This is important but boring,”* where the first adjective is considered positive in some sense, while the second one is negative. In addition, Expectation Violation MCs in conjunction with Temporal MCs (Halliday & Hasan 1976), signal a violation of timing expectations, e.g., *“Today we shall do X, but first let us revise Y.”*
- ii. Focal — Signal the focus or context of a message [*“Now, about his sister, ... ,” “By the way”*].
- iii. Implementational — Signal the informational status of a message, i.e., whether it contains new information or information to be verified against existing knowledge [*“Let us consider a new topic,” “As I said before”*].
- iv. Estimational — Indicate the length and/or difficulty of a message [*“Here comes the tricky part,” “The proof is rather long”*].
- v. Interpretive — Advise which “parsing mechanism” (style) should be used to interpret a message [*“If you pardon the expression”* — colloquial, *“Technically speaking”* — technical].
- vi. Affective[†] — Advise of the affective impact associated with a message [*“Unfortunately”*].

Focal, Implementational, Estimational and Interpretive MCs advise a listener of the value of an attribute, and are generated when this value causes a change in the status of an attribute expectation. Expectative MCs, on the other hand, advertise the type of change which has taken place in the status of a belief or expectation, e.g., violation or realization, without specifying the cause of the change.

The Basic Mechanism

In this section, we present algorithm *Determine-MC* which generates codes representing requirements for MCs with respect to attribute expectations[‡]. This mechanism is basically domain independent, however, it assumes a hierarchical discourse and domain structure. A domain hierarchy is also used to model a listener's knowledge, where each node contains an *Own Attribute List* and a *Descendant Attribute List*. The Own List contains the values of the length, difficulty, style and timing attributes of a node, and the Descendant List contains the focus, informational status and Own List attributes of an expected successor of this node. *Determine-MC* receives two arguments: a message sequence *M* and a starting node *N*, and assumes that the message sequence is compatible with a relevant portion of the hierarchy representing the listener's knowledge base. The activation of the algorithm starts at a

[†] The term *affects* is used throughout this paper in the sense used in [Dyer 1982], i.e., meaning “emotions.”

[‡] A similar algorithm is used to determine requirements for MCs with respect to content expectations, except that line 6, line 10, and the second disjunct in line 4 are omitted, since they are not applicable to content expectations.

ZUKERMAN

node arrived at by means of default expectations with respect to a given context, and the first message in the sequence corresponds to one of the children of this node.

Procedure Determine-MC (M, N)

```

1  If M is empty Then return nil
2  Create a list called MC-list which is initially set to head(M)
3  For each attribute aj in N do
4    If { [ head(M) causes a change in the status an expectation for aj ] or
        [ There are conflicting expectations with respect to aj ] }
        Then propose an MC-code and add it to MC-list
5    end
6  If { [ MC-list = head(M) ] and [ The information in head(M) was expected ] }
    Then let N' = N ; omit head(M)
7  Else
8    For each item in MC-list
9      Perform direct and indirect inferences with respect to expectations
10     Activate EADs where applicable
11     end
12    Let N' be the node put in focus by head(M)
13  Append (MC-list, Determine-MC (tail(M), N'))
14  end

```

To determine whether a message causes a change in the status of an expectation (line 4), we first consider the expectations in node N pertaining to the focus and informational status of this message. An active focus expectation in node N points to one of its children, and an expected informational status may be either *new*, *existing* or *repeated*. These expectations determine the need for a Focal and an Implementational MC, respectively. We now proceed to the node put in focus by the current message, denoted N' , in order to ascertain the need for additional MCs. We determine whether a change in the status of an attribute expectation has occurred by taking into consideration the Descendant List of N , the Own List of N' , and the information in the message. If N' has no Own List (e.g., it is a new node), or its Own List agrees with the Descendant List of N (e.g., it realizes a focus expectation), then the message attributes are compared against the Descendant List of N . Otherwise, if the value of attribute a_j in N disagrees with the value of a_j in N' , then an expectation violation has taken place even prior to considering the corresponding attribute in the message. Once a change in the status of an expectation has been detected, or the presence of conflicting active expectations is ascertained, the system proposes MCs where necessary (see next section). Finally, if all the information in a message is expected and no expectation realization or confirmation MCs were proposed (line 6), the entire message may be omitted. However, in this case, one must ensure that the listener's active expectations are sufficient to overcome the lack of continuity in the hierarchy traversal. (The issues related to this case are currently being investigated.)

Proposing MCs

The procedure which determines the need for MCs caters for three subgoals maintained by a speaker in a knowledge acquisition setting: (1) Transfer knowledge, (2) Enhance his/her own credibility, and (3) Foster a positive attitude in the listener. These subgoals may be violated if, in the absence of an MC, the user is likely not to notice a change in the status of an expectation, or to experience negative affects, such as confusion or loss of confidence in the system. The MCs proposed to prevent these undesirable effects depend on the nature of the change in the status of an expectation or belief (violation, realization or confirmation), the type of the expectation or belief (content or attribute), its explicitness, the level of certainty associated with it, and its affective impact.

A speaker's failure to generate an MC which signals a change in the status of a listener's expectations or beliefs may be interpreted by the listener as an indication that the speaker has either failed to notice this change, or that s/he does not consider this change worth mentioning (it may be either too obvious or too inconsequential). A listener's choice of interpretation and subsequent reaction depend both on his/her opinion regarding his/her knowledge status relative to the speaker's, and on his/her evaluation of the significance of the change in question. For example, consider the text "We factor out x hoping to find a common factor. *However*, we get $x(x+1)+7$," where the second sentence, stating the result of the factoring out operation, violates the explicit content expectation activated by the first sentence. If the Expectation Violation MC is omitted, one of the following events may take place: (1) a novice may not even notice that the expectation for a common factor has been violated, (2) a listener who does not feel confident in the subject matter might notice the expectation violation, but may conclude s/he is probably mistaken, since the speaker hasn't mentioned it, (3) a listener who considers him/herself familiar with the subject may lose respect for a speaker who hasn't mentioned (and hence probably hasn't noticed) the expectation violation, or (4) if both the listener and the speaker are proficient mathematicians, the expectation violation will be accepted without further ado. A similar analysis may be performed for the realization of content expectations. For instance, in the text "This procedure is polynomial. *Indeed*, it is used in many algorithms," the second sentence realizes an implicit expectation resulting from the inference that polynomial procedures are efficient and hence are used often. If the Expectation Realization MC is omitted, an inexperienced listener may completely fail to infer the relationship between the two sentences, a more experienced listener may be troubled by the fact that the speaker has neglected to point out this relationship, and an expert may simply carry on after making the appropriate inferences.

In principle, a similar scenario may take place for attribute expectations, where Focal, Implementational, Estimational and Interpretive MCs are generated to inform a listener of the cause of a change in the status of an expectation or a belief with respect to an attribute. However, since people are generally relatively competent with respect to discourse attributes and believe that most other people are equally competent in this area, case (4) occurs most frequently, obviating the need for Expectative MCs, and cases (1) and (2), on the other hand, are extremely rare. Case (3) may occur only due to the violation of a strong attribute expectation, e.g., "Today we shall do X. *On second thought*, let's do Y" (explicit focus expectation) and "*Contrary to custom*, we shall start with dessert" (implicit focus expectation).

Finally, Expectation Realization MCs are generated in additional cases which are not accounted for by this analysis, namely they directly contribute to the satisfaction of a speaker's subgoals of enhancing his/her credibility and fostering a positive attitude in the listener, by allocating credit to the speaker or the listener, respectively. The former is illustrated by "As I expected, Bush won the election," and the latter by "As you may know, this is a linear equation." In particular, if the Expectative MC is omitted from the second text, a listener who did possess the belief in question may resent a speaker who is underestimating his/her knowledge.

These arguments are operationalized as follows:

Expectation Violation MCs are generated to signal the violation of (1) beliefs or expectations with respect to the content of a forthcoming message, and (2) strong attribute expectations. **Expectation Realization MCs** are generated to indicate the realization or confirmation of (1) expectations or beliefs a listener is likely to entertain both with respect to content and attributes, and (2) weak content expectations or beliefs, where the weakness of a belief or an expectation may stem from its uncertainty or from the inferences which must be performed to activate it.

Focal, Implementational, Estimational and Interpretive MCs are generated to advertise the violation of attribute expectations, such as changes in the topic of conversation or the difficulty of the problems being considered, the realization of weak attribute expectations, and the resolution of conflicting attribute expectations.

Affective MCs are generated to signal the inactivation or confirmation of content and attribute expectations due to an event which carries an affective impact for the speaker or the listener and is associated with some uncertainty, e.g., "There is a way to solve this problem. *Unfortunately*, it is quite difficult," where the affective impact of an event is obtained from models of the speaker and the listener.

The explicitness of an English referent representing an Expectative MC depends on the number of concurrently active expectations and on the distance (in time and number of messages) between the last reference to an expectation and an inactivating or confirming result, e.g., whereas an implicit Expectative MC is suitable for an expectation currently in context, a more distant expectation may require a more explicit reference, such as "We were hoping to obtain a product of factors, however, we got" The lexical choices made for MCs which pertain to discourse related attributes depend on the type of the message in question and on the conditions surrounding its generation. For instance, an Implementational MC indicative of verification would be "revise" or "return" for a topic, and "too" or "again" for an algebraic operation.

We conclude this section with a trace of our mechanism which accounts for the output generated for the partial input in figure 1.

1	(topic <i>quadratic</i>)	} Let us go on with the topic of quadratic equations.
2	(equation $(x-3)^2-4(x-3)-7=0$)	
	(# alt = 3)	} An equation follows:
		} $(x-3)^2 - 4(x-3) - 7 = 0$
		} There are three ways of solving it.
3	(alternative 1)	} The first alternative consists of the following operations:
	{Solutions to the equation}	...

Fig. 1: Partial Input to FIGMENT II and its Corresponding Output[†]

Since our domain is algebra, the Default demon activates an initial expectation for the focus to be Math → Algebra. Now, the topic of quadratic equations has been discussed in the previous session, causing a Diversity demon to activate an expectation for a new topic, and a Default demon to activate a weak expectation for the continuation of quadratic equations based on length considerations, i.e.,

[†] At present, the input to FIGMENT II is hand-coded, however work is currently in progress on a complete interface between FIGMENT II and a system which produces human oriented solutions to algebraic equations (Oliver & Zukerman 1988).

ZUKERMAN

one expects to discuss a topic for more than one session. These expectations yield conflicts with respect to the focus and the informational status entries in the Descendant List of the Algebra node, which demand the generation of a Focal MC and a verification Implementational MC, e.g., "Let us continue with the topic of quadratic equations." The Quadratic Equations node is now visited, and any inferred expectations are incorporated into its Own List: a direct inference results in the activation of an explicit timing expectation for this topic to be discussed immediately, and the firing of a Default demon yields an implicit length expectation for this topic to be discussed for a default time period. Since a variety of subtopics of quadratic equations may have been discussed in the previous session, there isn't a strong expectation for any of them. Hence, a Focal MC, such as "Here is an equation" is required. Notice, however, that the type of activity, namely problem solving, does not need to be mentioned, since, in this case, the mere presentation of an equation triggers a strong expectation for problem solving. Since a Default length expectation for one solution alternative is active in the Own List of the node corresponding to the equation and the number of presented solution alternatives violates this expectation, an Estimational MC is required, e.g., "There are three ways of solving this equation." This MC, however, points to all three alternatives, thereby weakening the focus expectation in the Descendant List of the Equation node. Hence, a Focal MC for each alternative is called for.

RELATED RESEARCH

The research reported in this paper builds on the mechanism presented in [Zukerman & Pearl 1986] which generates a subclass of MCs that signal a discrepancy between a given message and "normal" expectations. The current mechanism is able to provide a more uniform account for the generation of a broader class of MCs by dynamically drawing inferences from the state of the discourse to anticipate a listener's expectations. In addition, this mechanism accounts both for the presence and the absence of MCs in situations where the previous mechanism failed to provide a competent explanation.

Taxonomies of a subclass of MCs, namely conjunctive expressions, appear in [Halliday & Hasan 1976], [Longacre 1976] and others. A functional taxonomy encompassing additional Meta Comments appears in [Zukerman & Pearl 1986]. Theoretical work on Interpretive MCs appears in [Sigurd 1986], and several researchers (Cohen 1987, Reichman-Adar 1984, Grosz & Sidner 1985, McCoy & Cheng 1988) have presented mechanisms which view Focal Meta Comments as indicators of non-default traversals of a structure representative of the discourse. Finally, a detailed discussion on expectations can be found in [Ortony & Partridge 1987].

CONCLUSIONS

In this paper, we have introduced a mechanism for the generation of Meta Comments based on the tenet that they signal the inactivation or confirmation of a listener's beliefs and expectations. We have presented a simple knowledge representation scheme which supports this mechanism with respect to attribute and content expectations. To support this mechanism with respect to beliefs, a representation of a listener's beliefs resulting from previous discourse is required (Zukerman & Cheong 1988). In addition, we have introduced the concept of Expectation Activation Demons which emulate human behaviour in the activation of attribute expectations. We have incorporated our model in a text generation facility, demonstrating its use as a tool for the generation of fluent and cogent text.

REFERENCES

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

ZUKERMAN

- Cohen, R. (1987), Interpreting Clues in Conjunction with Processing Restrictions and Arguments and Discourse. In *AAAI-87 Conference Proceedings*, Seattle, August 1987.
- Dyer, M.G. (1982), *In-depth Understanding: A Computer Model of Integrated Processing for Narrative Comprehension*. Doctoral Dissertation, Department of Computer Science, Yale University, New Haven, Connecticut.
- Grosz, B.J. & Sidner, C.L. (1985), Discourse Structure and the Proper Treatment of Interruptions. In *IJCAI-85 Conference Proceedings*, Los Angeles, August 1985.
- Halliday, M.A.K. & Hasan, R. (1976), *Cohesion in English*. Layman Press, London.
- Hovy, E.H. (1987), Pragmatics and Natural Language Generation. Technical Report, Information Sciences Institute, Los Angeles, July 1987.
- Kukich, K. (1983), *Knowledge-Based Report Generation: A Knowledge-Engineering Approach to Natural Language Report Generation*. Doctoral Dissertation, The Interdisciplinary Department of Information Science, University of Pittsburgh, Pennsylvania.
- Longacre, R.E. (1976), *An Anatomy of Speech Notions*, Peter de Ridder Press, Publications in Tagmemics No. 3.
- McCoy, K.F. & Cheng, J. (1988), Focus of Attention: Constraining What Can Be Said Next. *Unpublished Manuscript*.
- McKeown, K.R. (1985), Discourse Strategies for Generating Natural Language Text. In *Artificial Intelligence 27*, pp. 1-41.
- Oliver, J. & Zukerman, I. (1988), DISSOLVE: A System for the Generation of Human-Oriented Solution to Algebraic Equations. In *Proceedings of AI'88*, Adelaide, November 1988.
- Ortony, A. & Partridge, D. (1987), Surprisingness and Expectation Failure: What is the Difference? In *IJCAI-10 Conference Proceedings*, Milan, August 1987.
- Paris, C.L. (1987), Combining Discourse Strategies to Generate Descriptions to Users along a Naive/Expert Spectrum. In *IJCAI-10 Conference Proceedings*, Milan, August 1987.
- Reichman-Adar, R. (1984), Extended Person-Machine Interface. In *Artificial Intelligence 22*, pp. 157-218.
- Riesbeck, C.K. (1982), Realistic Language Comprehension. In *Strategies for Natural Language Processing*, W.G. Lehnert and M.H. Ringle (Eds.), Lawrence Erlbaum Associates.
- Sigurd, B. (1986), Meta Comments in Text Generation. In G. Kempen (Ed.), *Natural Language Generation: New Results in Artificial Intelligence, Psychology and Linguistics*, pp. 453-461, Kluwer Academic Publishers.
- Swartout, W.R. (1982), XPLAIN: A System for Creating and Explaining Expert Consulting Programs, USC/Information Sciences Institute.
- Zukerman, I. (1986), *Computer Generation of Meta-technical Utterances in Tutoring Mathematics*. Doctoral Dissertation, Computer Science Department, University of California, Los Angeles.
- Zukerman, I. & Pearl, J. (1986), Comprehension-Driven Generation of Meta-technical Utterances in Math Tutoring. In *AAAI Conference Proceedings*, August 1986.
- Zukerman, I. & Cheong, Y.H. (1988), Impairment Invalidation: A Computational Model for the Generation of Rhetorical Devices. In *Proceedings of the International Computer Science Conference '88: Artificial Intelligence, Theory and Applications*, Hong Kong, December 1988.