Considering Explanation Failure during Content Planning

Ingrid Zukerman
Department of Computer Science
Monash University
Clayton, Victoria 3168, AUSTRALIA
ingrid@bruce.cs.monash.edu.au

Richard McConachy
Department of Computer Science
Monash University
Clayton, Victoria 3168, AUSTRALIA
ricky@bruce.cs.monash.edu.au

Abstract

Content planning systems generate explanations to achieve a communicative intent, often with respect to a particular audience. However, current research in content planning does not take into consideration the fact that an addressee may stop paying attention to an explanation because of boredom or cognitive overload. In this case, the generated explanation fails to achieve the communicative intent. In this paper, we present a computational representation of boredom and cognitive overload, and cast the problem of content planning as a constraint-based optimization problem. The objective function in this problem is a probabilistic function of a user's beliefs, and the constraints are restrictions placed on the boredom and cognitive overload the user can experience, and on the minimal level of expertise the user should achieve. We discuss two techniques for solving the optimization problem, and consider two types of constraints for addressing the expertise requirements. We also examine how variations in the populations to which constraints are applied affect the generated discourse.

Introduction

Current research in content planning focuses on planning discourse that achieves an intended communicative goal, often with respect to a particular audience, e.g., (Paris, 1988; Moore & Swartout, 1989; Cawsey, 1990; Zukerman & McConachy, 1993b). However, this research does not take into consideration the fact that an addressee may get bored or may be unable to follow the thread of the explanation, and stop listening to the speaker. In this case, regardless of how carefully the discourse has been planned, it will fail to achieve the intended communicative goal.

In this paper, we present a content planning mechanism which addresses this problem. Our mechanism generates a set of *Rhetorical Devices (RDs)* which achieves as much of the intended communicative goal as possible, while ensuring that the user does not experience boredom or cognitive overload, and that s/he achieves at least a minimum level of expertise. An RD consists of a rhetorical action, such as Assert or Instantiate, applied to a proposition. A set of RDs typically contains (1) RDs that convey the intended propositions, (2) subordinate sets of RDs that present prerequisite or referring information, and (3) RDs that address erroneous inferences. These RDs are then organized by the discourse structuring procedure described in (Zukerman & McConachy, 1993a), and rendered in English by the Functional Unification Grammar described in (Elhadad, 1992).

Table 1 illustrates two types of discourse generated by our mechanism to convey information about photosynthesis to a fairly weak student. The discourse on the left-hand-side contains all the information that the system deems necessary to achieve the intended communicative goal. In contrast, the discourse on the right-hand-side is considerably shorter, since the system anticipates that the student will get bored with the longer text, and hence decides to present only the RDs that convey the gist of the communicative goal.

In the following section, we discuss factors that cause explanation failure. Next, we specify the optimization problem, and describe two procedures that solve this problem: (1) full optimization, and (2) partial optimization combined with heuristics. We then analyze the effect of variations in constraints on the discourse generated by our system, and present concluding remarks.

Factors that cause explanation failure

An explanation fails to achieve an intended communicative goal when the addressee stops listening to it. Two factors that lead to this behaviour are boredom and cognitive overload. Our measures for anticipating whether the discourse being planned is likely to cause boredom or overload are based on the following observations.

Boredom due to excessive discourse length. Most people find many instructional texts inherently uninteresting. Hence they get bored if instructional discourse is too long. The length of discourse a user can tolerate depends on his/her ability and attitude, e.g., weak students are generally less able to cope with long discourse than stronger students.

Boredom due to unnecessary RDs. People get bored if discourse contains information that they already know or can infer easily. When a speaker addresses a heterogeneous audience, and aims the presentation at the average addressee, the addressees that are better than average are likely to experience this type of boredom.

Cognitive overload. People experience cognitive overload when they cannot integrate the pieces of an explanation to form a coherent body of knowledge [Just & Carpenter, 1987]. This can happen because they need to follow a lengthy argument or assimilate large amounts of background or clarifying information. Lengthy argu-

Table 1: Sample discourse with/without considering boredom

Without considering boredom	Considering boredom
Plants photosynthesize. This process consumes water, carbon-dioxide and nutrients, and produces glucose and oxygen. Plants use light, which is a form of energy, to photosynthesize. However, plants cannot use other forms of energy to photosynthesize. Plants contain chlorophyll, which is a green pigment that absorbs light. However, chlorophyll does not absorb green light. In addition to photosynthesizing, plants respire. Respiration consumes oxygen and glucose, and produces carbon dioxide and water. Plants photosynthesize more than they respire. Therefore, plants produce more oxygen than they consume.	Plants photosynthesize. This process produces oxygen. Plants also respire. This process consumes oxygen. Plants photosynthesize more than they respire. Therefore, plants produce more oxygen than they consume.

ments are often required when the addressee has to make large changes in belief; while large amounts of background or clarifying information are indicative of relatively small shifts with respect to several beliefs. Cognitive overload also takes place when an addressee has difficulty understanding parts of an explanation. However, this situation is identified in our system, and corrected by generating supporting, background or clarifying information, thus yielding the two above mentioned cases.

The prediction of boredom and cognitive overload requires information regarding the length of discourse, the amount of known or easily inferred information, and the maximum shift in belief an addressee can tolerate. In the current research, these thresholds are determined for stereotypical user models by testing the effect of different values on the discourse, while ensuring that they make sense relative to each other. For example, strong students can usually cope with a higher cognitive load than weak students.

Predicting boredom due to excessive length

We compare the length of the discourse with the maximum length which an addressee can presumably tolerate.

A good approximation of the length of a piece of discourse is simply the number of RDs in the discourse $|\{RD\}_{said}|$. Clearly, this approximation can sometimes be wrong, but it is not productive to realize a piece of discourse currently being planned just to measure its exact length. Thus, the requirement to avoid boredom due to excessive length is expressed by the following formula:

$$|\{RD\}_{said}| \leq TRD_{max}(M) \qquad (1)$$

where $TRD_{max}(M)$ is the maximum number of RDs a user who belongs to user model M can tolerate.

Predicting boredom due to unnecessary RDs

We determine the amount of planned discourse that is already known or can be easily inferred by the addressee, and compare it with the amount of such superfluous discourse that the addressee can presumably tolerate.

The superfluous part of the discourse is the difference between the planned discourse and what really needs to be said to achieve a communicative goal. Thus, the requirement to avoid boredom due to unnecessary RDs is expressed by the following formula:

$$|\{RD\}_{said} - \{RD\}_{reduced}(M)| \leq URD_{max}(M) \qquad (2)$$

where $\{RD\}_{said}$ is the set of RDs generated, $\{RD\}_{reduced}(M)$ is the most reduced version of $\{RD\}_{said}$ that can still convey the intended propositions to an addressee who belongs to a particular model M, and $URD_{max}(M)$ is the maximum number of unnecessary RDs a user who belongs to user model M can tolerate.

$\{RD\}_{reduced}(M)$ is obtained by removing RDs from $\{RD\}_{said}$ so long as the communicative goal is still achieved with respect to model M. When removing an RD, we also remove the RDs that depend only on this RD, i.e., the sets of RDs that convey prerequisite and referring information for this RD only, and the RDs that contradict erroneous inferences from this RD.

Predicting cognitive overload

We compare the total shift in belief required by the communicative goal with an addressee's conjectured ability to tolerate shifts in belief.

The total shift in belief required to achieve a communicative goal depends on an addressee's current beliefs and on the inferences s/he is likely to make. For instance, a user who has strong erroneous beliefs will need to make more adjustments than a user whose beliefs are close to the intended ones.

We distinguish between three types of propositions for the purpose of predicting cognitive overload: P – propositions that were previously unknown or correctly believed by the user; P' – propositions that were wrongly believed by the user and must now be contradicted; and $\hat{P}$ – propositions that were wrongly inferred by the user as a result of discourse planned to convey P or $\neg P'$ and must now be contradicted. The difficulty associated with the different types of shifts in belief is represented by F factors as follows. $F_P(M)$ reflects the amount of effort required to acquire new information, $F_{P'}(M)$ reflects the amount of effort required to reverse a previous belief, and $F_{\hat{P}}(M)$ reflects the amount of effort required to reverse a new inference ($F_P(M) < F_{\hat{P}}(M) < F_{P'}(M)$). These factors depend on the type of the user, e.g., a strong student usually has stronger convictions than a weak

Table 2: Propositional input that yields the sample texts

Propositions to be Conveyed	Intended Degree of Belief	Significance
*[plants do-action photosynthesis]	*BELIEVED*	CRITICAL
[photosynthesis consume water]	*RATHER BELIEVED*	MEDIUM
[photosynthesis consume carbon-dioxide]	*RATHER BELIEVED*	MEDIUM
[photosynthesis consume nutrients]	*BELIEVED*	HIGH
[photosynthesis produce glucose]	*RATHER BELIEVED*	MEDIUM
*[photosynthesis produce oxygen]	*BELIEVED*	CRITICAL
[plants do-action photosynthesis (use light)]	*BELIEVED*	HIGH
[plants contain chlorophyll]	*RATHER BELIEVED*	HIGH
*[plants do-action respiration]	*BELIEVED*	CRITICAL
*[respiration consume oxygen]	*BELIEVED*	CRITICAL
[respiration consume glucose]	*RATHER BELIEVED*	MEDIUM
[respiration produce carbon-dioxide]	*RATHER BELIEVED*	MEDIUM
[respiration produce water]	*RATHER BELIEVED*	MEDIUM
*[plants do-action photosynthesis(X)] *[plants do-action respiration(Y)] $\}$ * $[X > Y]$	*RATHER BELIEVED*	HIGH
*[plants produce oxygen(Z)] *[plants consume oxygen(W)] $\}$ * $[Z > W]$	*RATHER BELIEVED*	HIGH

student, and hence will have more difficulty reversing a belief.

The following formula expresses the total weighted shift in belief experienced by a user who belongs to model M when attempting to achieve an intended degree of belief with respect to a set of propositions.

$$T_{SHIFT}(M) = \sum_{p \in P} f_{SHIFT_M}(p) \cdot F_P(M) + \sum_{p \in P'} f_{SHIFT_M}(p) \cdot F_{P'}(M) + \sum_{p \in \hat{P}} f_{SHIFT_M}(p) \cdot F_{\hat{P}}(M)$$

where $f_{SHIFT_M}(p)$ represents the contribution of proposition p to $T_{SHIFT}(M)$. This contribution is the absolute value of the difference between the actual and the previous belief in p for a user who belongs to model M [1].

$$f_{SHIFT_M}(p) = |bel_{act_M}(p) - bel_{old_M}(p)|$$

Thus, the requirement to avoid a total shift in belief which results in cognitive overload is expressed by the following formula:

$$T_{SHIFT}(M) \leq belshift_{max}(M) \qquad (3)$$

where $belshift_{max}(M)$ is the maximum shift in belief a user who belongs to model M can tolerate.

The Content Planner

Our mechanism uses a constraint-based optimization procedure whose objective is to maximize a user's belief with respect to a set of intended propositions, and whose constraints are restrictions placed on the boredom and cognitive overload the user may experience, and on the minimum level of expertise the user should achieve.

Our system receives two types of input: propositional and user-model related. The propositional input contains (1) a set of propositions to be conveyed; (2) the degree of belief the user is expected to achieve with respect to each proposition; and (3) the significance of each proposition, i.e., how important it is that the user believes it. These last two inputs are often correlated. However, using separate measures for these inputs allows us to model cases that differ from the norm. For example, when conveying background information for an intended proposition it is crucial that a user acquire at least a passing acquaintance with the information in question. Table 2 shows the propositional input from which the texts in Table 1 are generated. The text on the left-hand-side of Table 1 conveys all the input propositions, while the text on the right-hand-side conveys only the propositions marked with an asterisk (*).

The user-model related input is a space of user models accompanied by a probability distribution. The probability distribution may be interpreted either as the system's uncertainty regarding which model(s) a particular user belongs to, or as the percentage of a group of addressees that belongs to each of these models. The models are used to determine the information that needs to be presented to achieve a given communicative goal with respect to a particular audience. Each model represents the beliefs, inference patterns and attitude of a particular type of user. The attitude models a user's ability to understand abstract information, his/her confidence in his/her inferences, and the length of discourse, cognitive overload and amount of unnecessary RDs s/he is likely to tolerate. In the current implementation we maintain five stereotypical user models: ***excellent***, ***good***, ***average***,

[1] The 'actual' degree of belief is conjectured by means of a function which simulates a user's change in belief as a result of a piece of discourse. This function depends on the user's ability and on the complexity and abstractness of the information (Zukerman & McConachy, 1993b).

Table 3: Set of RD that yields the unconstrained sample text

Assert[plants do-action photosynthesis]	Assert[plants contain chlorophyll]
Assert[photosynthesis consume water]	◇Assert[chlorophyll isa green-pigment]
Assert[photosynthesis consume carbon-dioxide]	◇Assert[chlorophyll absorbs light]
Assert[photosynthesis consume nutrients]	◇Negate[chlorophyll absorbs green-light]
Assert[photosynthesis produce glucose]	Assert[plants do-action respiration]
Assert[photosynthesis produce oxygen]	◇Mention[plants do-action photosynthesis]
Assert[plants do-action photosynthesis (use light)]	Assert[respiration consume oxygen]
◇Assert[light isa form-of-energy]	Assert[respiration consume glucose]
◇Negate[plants do-action photosynthesis (use other-form-of-energy)]	Assert[respiration produce carbon-dioxide]
	Assert[respiration produce water]
Compare[{plants do-action photosynthesis} > {plants do-action respiration}]	
Compare[{plants produce oxygen} > {plants consume oxygen}]	

mediocre and *weak* (a detailed description of these models appears in (Zukerman & McConachy, 1993b)).

The output of the content planner is a set of RDs which are related to each other by means of discourse relations such as prerequisite and causality (Mann & Thompson, 1987). Table 3 contains the set of RDs which yields the text on the left-hand-side in Table 1 (without the discourse relations). The RDs marked with a diamond (◇) convey background information and contradict erroneous inferences.

Specification of the Optimization Process

The objective of the optimization process is to plan discourse that minimizes the distance between the actual and the intended degree of belief with respect to a list of intended propositions without violating the boredom, overload and minimum-expertise constraints. The belief objective must be achieved probabilistically with respect to all user models. Further, we emulate behaviour whereby speakers make sure that important propositions are conveyed, while placing less emphasis on less important propositions. To this effect, we take into consideration both the significance of the intended propositions and the degree to which these propositions are to be believed upon completion of the discourse. This yields the following objective function for the optimization process:

$$\min\{\sum_M\{\sum_p f_{BEL}(p)Sig(p)\}Prob(M)\} \qquad (4)$$

where

$$f_{BEL_M}(p) = \begin{cases} 0 & \text{if } |bel_{act_M}(p) \geq bel_{int}(p)| \text{ and } sign(bel_{act_M}(p)) = sign(bel_{int}(p)) \\ |bel_{act_M}(p) - bel_{int}(p)| & \text{otherwise.} \end{cases}$$

$Sig(p)$ is the significance or importance of a proposition, and $f_{BEL}(p)$ represents the contribution of proposition p to the objective function. This contribution is 0 if $bel_{act_M}(p)$ exceeds or equals the intended belief in p, $bel_{int}(p)$, for a user who belongs to model M. Otherwise, it is the absolute value of the difference between the intended and the actual belief in p.

The boredom and overload constraints are expressed by Equations (1-3). However, in order to take into consideration the system's uncertainty regarding which model a user belongs to, we moderate the thresholds by a function which relaxes the constraints as the probability that the user belongs to a particular model decreases. The following formulation specifies one constraint for each type of explanation failure and each user model:

For $i \in$ {*excellent, good, average, mediocre, weak* }

Boredom (length)

$$|\{RD\}_{said}| \leq TRD_{max}(M_i) \cdot f_{Prob}(M_i) \qquad (5)$$

Boredom (unnecessary RDs)

$$|\{RD\}_{said} - \{RD\}_{reduced}(M_i)| \leq URD_{max}(M_i) \cdot f_{Prob}(M_i) \qquad (6)$$

Overload

$$T_{SHIFT}(M_i) \leq belshift_{max}(M_i) \cdot f_{Prob}(M_i) \qquad (7)$$

In addition, we propose two alternative formulations for expertise-related constraints:

$$\frac{1}{|\{p\}|}\sum_M\{\sum_p f_{BEL_M}(p)Sig(p)\}Prob(M) < Thr \qquad (8)$$

For $i \in$ {*excellent, good, average, mediocre, weak* }

$$\sum_p bel_{act_{M_i}}(p)Sig(p) > Prct_{M_i}\sum_p bel_{int}(p)Sig(p) \qquad (9)$$

Equation 8 stipulates that the average shortfall from the intended level of expertise weighted across all user models should not exceed a certain threshold (the shortfall is averaged over the number of propositions to be conveyed). This constraint, which is placed on the objective function, ensures that an adequate level of expertise is achieved on average for each proposition over all the user models. In contrast, Equation 9 demands that a user who belongs to a particular model attain a certain percentage of the intended level of expertise. For example, a weak student may be expected to attain at least 50% of the intended expertise, while an excellent student may be expected to attain at least 90%.

The above constraints are not necessarily applied to all user models with non-zero probability. Given a (possibly 0) top margin and bottom margin, the system drops from consideration the models at the tail of the distribution whose probabilities fall within these margins. For instance, an input such as (10%, 20%) means that the constraints are not applied with respect to the model(s) whose probabilities fall within the top 10% or the bottom 20% of the user-model distribution. In addition, we have used Equation 9 in two different ways: (1) applied to all user models that are being considered, and (2) applied to the weakest user model under consideration.

The above formulation yields a non-linear integer optimization problem even without the constraints. Hence, a weak search method is applied. We have implemented two methods to solve this problem: (1) a full optimization (Zukerman & McConachy, 1995), and (2) a partial optimization combined with heuristics. If these methods cannot find a solution that satisfies all the constraints, they relax the communicative goal, i.e., they convey some propositions to a lesser extent than originally specified and/or give up conveying some propositions altogether. This approach is suitable for a situation where time is running short, and the information provider is willing to forego the conveyance of non-essential information so long as crucial information is conveyed. An alternative approach which consists of relaxing the single-discourse requirement, i.e., conveying the same amount of information in separate chunks of discourse, is described in (Zukerman & McConachy, 1995).

Full optimization

This procedure first determines *minimally sufficient* sets of RDs which convey all the propositions for each user model, where a set of RDs is minimally sufficient if the removal of any RD causes the set to stop conveying the intended information. It then iteratively selects the set of RDs with the best objective function among the sets of RDs which satisfy all the constraints, and generates RDs that convey prerequisite and referring information for the selected set of RDs. The resulting set of RDs, i.e., the selected set plus the RDs which convey its referring and prerequisite information, is added to the pool of candidate sets, and its constraints are re-calculated. The optimization process terminates when it finds a set of RDs which satisfies all the constraints and requires no additional prerequisite or referring information.

If all the candidate sets of RDs violate one or more constraints, the communicative goal is relaxed as follows. The system successively removes one RD from each set of RDs that satisfies the expertise constraint(s), and inspects the effect of each removal on the constraint(s) and the objective function. During this process, when removing an RD we also remove the RDs that depend only on this RD (as when generating $\{RD\}_{reduced}$ while predicting boredom due to unnecessary RDs). Each reduced set of RDs which results from this process satisfies all the constraints while yielding an objective function whose value is worse than before.

Partial optimization

This procedure first applies the mechanism described in (Zukerman & McConachy, 1994) to generate the optimal set of RDs for each of the user models being considered. The sets of RDs are then sorted with the value of the objective function as the primary sorting key, and the number of constraints violated as the secondary key. The highest-ranked set that violates no constraints is then selected.

If all the sets of RDs violate one or more constraints, the communicative goal is relaxed as follows. The system selects the top-ranked set of RDs, and applies rules which take into consideration the significance and the intended shift in belief of the propositions to be conveyed in order to select a proposition which may be conveyed to a lesser extent (or not at all). These rules select first propositions of low significance, next propositions of medium significance which require a low or medium shift in belief, and so on. One of the RDs that conveys the selected proposition is then removed, and the constraints and objective function are re-calculated. If the resulting set of RDs violates no constraints, it is selected for presentation. Otherwise, the set of RDs is re-ranked according to the value of its objective function and the number of constraints it violates, and the process is repeated.

Results

The system was run with several combinations of the following parameters: (1) the two optimization procedures; (2) the two constraints for minimum expertise, where Equation 9 was applied in two modes: to all user models and only to the weakest user model; and (3) the two types of target populations, viz all the user models and only user models that fall inside specified margins. The application domains included technical areas such as nuclear fission, chemistry and biology.

When boredom constraints are turned off there is no penalty for excessive length or unnecessary RDs, hence the generated text includes examples, background information and elaborations to ensure that the material is conveyed (left-hand-side of Table 1). When constraints pertaining to boredom due to length are activated, RDs that convey propositions of lower significance tend to be omitted first (right-hand-side of Table 1). If a proposition with a higher significance requires many RDs, then these RDs become good candidates for omission.

When overload constraints are activated, propositions that require a large shift in belief are removed. This is illustrated in the example in Table 4, where the RD that conveys proposition 1 and its dependent, the RD that conveys proposition 2, are removed. In contrast, when constraints pertaining to boredom due to length are activated for this example, propositions 3 and 4 are omitted (proposition 6 remains because it is linked to proposition 5).

When both types of boredom constraints are activated, if the probabilities of the user models are evenly distributed, the only way to satisfy both sets of constraints is to convey very little information, yielding an objective function with a high value. In this case, following accepted teaching practices, the system relaxes the unnecessary-RDs constraints, i.e., it gives a higher priority to the requirements of the weaker user models.

The full optimization procedure takes between 30-60

Table 4: Sample discourse considering overload versus boredom due to length

Propositions to be Conveyed	Intended Degree of Belief	Significance
1. Network covalent substances sublime,	*BELIEVED*	HIGH
2. which is going from a solid state directly to vapor.	*RATHER BELIEVED*	HIGH
3. Ionic substances melt.	*RATHER BELIEVED*	MEDIUM
4. Metallic substances melt.	*RATHER BELIEVED*	MEDIUM
5. Metallic substances are flexible,	*RATHER BELIEVED*	HIGH
6. but ionic substances are not flexible.	*RATHER BELIEVED*	MEDIUM

seconds of CPU time on a SPARCstation2. The partial optimization combined with the heuristic function cuts the processing time significantly (by 1/4 – 2/3 depending on the number of violated constraints and the relationship between the intended propositions). In general, the full optimizer conveys more propositions (at least partially) than the partial optimizer, and yields a better objective function. In contrast, the partial optimizer generally achieves larger belief shifts for the propositions that are not removed. There is a marked difference between the discourse generated by the partial and the full optimizer when the content planner initially generates a small set of RDs to convey a few propositions that are strongly related to each other (where the RDs generated to convey one proposition affect the others). In this case, the partial optimizer relaxes the communicative goal by removing one of the RDs from the set of RDs, while the full optimizer often replaces the entire set of RDs with a different, smaller set.

The two formulas representing expertise-related constraints affect the system's output as follows. If Equation 9 is used with respect to models that are not strongly represented in the user population, the output will be markedly different from the output obtained when Equation 8 is used. This happens because Equation 9 forces the system to address the needs of these models, while Equation 8 largely ignores these models owing to their small relative probability.

Finally, ignoring portions of the population that fall below the bottom margin yields more concise text, since less explanations need to be presented, while ignoring the portions that fall above the top margin allows the system to generate more RDs without violating the unnecessary-RDs constraints.

Conclusion

We have offered a computational definition of three causes of explanation failure: cognitive overload, boredom due to excessive discourse length and boredom due to unnecessary RDs. We have cast content planning as a constraint-based optimization process which takes into account a speaker's uncertainty regarding the user model to which an addressee belongs. In this process, the constraints represent requirements placed on the addressee's boredom and overload, and on the level of expertise s/he is expected to attain, and the objective function is a probabilistic function of the extent to which the communicative goal has been achieved. We have discussed two procedures for solving this problem in combination with two types of expertise-related constraints, and we have considered the effect of ignoring segments of the user population that are at the tail of the distribution.

References

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

Elhadad, M. (1992). FUG: The Universal Unifier User Manual Version 5.0 (Technical Report). Columbia University, New York, New York.

Just, M.A. & Carpenter, P.A. (1987). *The Psychology of Reading and Language Comprehension*, Allyn and Bacon, Inc.

Mann, W.C. & Thompson, S.A. (1987). Rhetorical Structure Theory: A Theory of Text Organization (Technical Report ISI/RS-87-190). Information Sciences Institute, Los Angeles, California.

Moore, J.D. & Swartout, W.R. (1989). A Reactive Approach to Explanation. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence* (pp. 1504-1510), Morgan Kaufmann Publishers.

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

Zukerman, I. & McConachy, R.S. (1993a). An Optimizing Method for Structuring Inferentially Linked Discourse. In *AAAI-93 Proceedings – the National Conference on Artificial Intelligence* (pp. 202-207), AAAI Press.

Zukerman, I. & McConachy, R.S. (1993b). Consulting a User Model to Address a User's Inferences during Content Planning. In *User Modeling and User Adapted Interaction* 3(2), 155-185.

Zukerman, I. & McConachy, R.S. (1994). Being Concise versus Being Shallow: Two Competing Discourse Planning Paradigms. In *ECAI94 Proceedings – The European Conference on Artificial Intelligence* (pp. 515-519), John Wiley & Sons.

Zukerman, I. & McConachy, R.S. (1995). Generating Discourse across Several User Models: Maximizing Belief while Avoiding Boredom and Overload. To appear in *IJCAI-95 Proceedings – the International Joint Conference on Artificial Intelligence*, Morgan Kaufmann Publishers.