Taking Initiative in Response Generation

Nancy Green
School of Computer Science and Department of Philosophy
Carnegie Mellon University
Pittsburgh, PA 15231-3891, USA
(Internet: nancy.green@cs.cmu.edu)

Abstract

We present an approach to taking the initiative to provide unrequested but relevant information in responses to Yes-No questions. In our model, stimulus conditions represent types of discourse contexts in which a speaker would be motivated to provide unrequested information. Stimulus conditions can be thought of as situational triggers which give rise to new speaker goals, and which are the compiled result of "deeper" planning based upon intentional goals. Stimulus conditions are used as part of a discourse-plan based approach to response generation which enables the generation of coherent multiple-sentence responses, including indirect answers.

Introduction

We present our approach to taking the initiative to provide unrequested but important relevant information in responses to Yes-No questions. This approach was developed as part of a computational model for generating and interpreting indirect answers to Yes-No questions (Green 1994). (By an *indirect answer* we mean a response which is intended to convey the answer but does not necessarily semantically entail the answer.) In that model, an indirect answer is generated in two phases. First, during the *content planning* phase, the system (modeling a speaker) decides what information to include, if any, in addition to the requested information. During the second phase, the system determines if the direct answer needs to be given explicitly or if it is inferrable from the additional information. In the former case, the system generates a response containing the direct answer as well as the additional information; in the latter case, the additional information is given instead of a direct answer. In this paper, we focus on the content planning phase, presenting the mechanisms used in our model for taking the initiative to provide unrequested but relevant information.

To give an example,[1] in (1) the responder (labelled R) has provided an explanation of why R will not have enough time although the questioner (labelled Q) has not explicitly asked for that information.

1 Q: `What about coming here on the way or doesn't that give you enough time?`
R: `Well no I'm supervising here`

According to one study of questions and responses in English (Stenström 1984), 85% of responses to Yes-No questions containing a *direct answer* were accompanied by unrequested but relevant information. Thus, this type of initiative – whether manifested as an indirect answer or in the occurrence of unrequested information accompanying a direct answer – is common in dialogue.

In the rest of this paper, we first describe some general types of contexts in which a responder would be motivated to provide this type of information. Following that, we describe how control of this type of initiative is modeled in our system. Finally, we compare our approach to previous work.

Stimulus Conditions

In this section, we describe five general types of situations, which we refer to as *stimulus conditions*, in which a responder would be motivated to provide unrequested but relevant information in his or her response to a question. (The method by which the unrequested information is selected is outlined in the following section.) For each stimulus condition, we describe how a computational system can detect that the stimulus condition holds in the current discourse context.

Our methodology for identifying stimulus conditions was to survey linguistic studies as well as to analyze the possible motivation of the speaker in the examples in our corpus of responses. While stimulus conditions may be derivative of "deeper" principles of cooperation (Grice 1975) or politeness (Brown & Levinson 1978), they provide a level of compiled knowledge about relevant features of the discourse context and user model which reduces the amount of on-line reasoning required by the system. However, with each stimulus condition, we also provide an informal description of the deeper principle of which the stimulus condition could be a manifestation.

[1] (Levinson 1983), example (55).

Since principles of politeness may be unfamiliar to the reader, we first provide a brief overview. Brown and Levinson (Brown & Levinson 1978) claim that politeness strategies account for many uses of language. According to Brown and Levinson, certain communicative acts are intrinsically *face-threatening acts* (FTAs). That is, doing an FTA is likely to injure some conversational participant's *face*, or public self-image. For example, orders and requests threaten the recipient's *negative face*, "the want ... that his actions be unimpeded by others" (p. 67). On the other hand, disagreement or bearing "bad news" threatens the speaker's *positive face*, the want to be looked upon favorably by others.

The following strategies given by Brown and Levinson are relevant to this work.

1. perform the FTA with redressive action, i.e., in a manner which indicates that no face threat is intended, using *positive politeness* strategies (strategies which increase the hearer's positive face). Such strategies include
   - *Strategy 1:* attending to the hearer's interests or needs
   - *Strategy 6:* avoiding disagreement
2. perform the FTA with redressive action, using *negative politeness* strategies (strategies for increasing negative face). These include
   - *Strategy 6:* giving an excuse or an apology

Returning to the main topic of this section, example (1) provides an illustration of the **Explanation-Indicated** stimulus condition, which may be glossed as follows:

Explanation-indicated: R suspects that Q suspects that a proposition p holds (when in fact, p does not).

(We use *suspect* as a gloss for holding a belief with a confidence level which is less than certainty.) The system could determine that this condition holds based upon its user model, the dialogue history, or syntactic clues. In (1), for example, the form of the question provides weak evidence that Q suspects that the questioned proposition is true (i.e., that coming here would give R enough time). This stimulus condition can be justified on grounds of cooperation, since simply informing the questioner that p does not hold may in turn motivate Q to ask a follow-up question of *why not?*.

Example (1) also provides an illustration of the **Excuse-Indicated** stimulus condition.

Excuse-indicated: R suspects or believes that Q's question is a *pre-request* of a request.

To give some background on the manner in which we have characterized this condition, a Yes-No question may be interpreted as a *pre-request* for another act, e.g., when used to determine if a precondition of another action holds (Levinson 1983). (In the speech act literature, making a pre-request in lieu of the main request is referred to as making an *indirect request*.) In (1), the question *Doesn't that give you enough time?* functions as a pre-request for the main request *What about coming here on the way?* (which itself may be an indirect form of invitation). Techniques for interpreting indirect speech acts (Perrault & Allen 1980; Hinkelman 1989) can be used to determine when **Excuse-indicated** obtains. This stimulus condition can be justified on the basis of politeness. A negative answer to a Yes-No question used as a pre-request may be interpreted as a refusal of the main request. To soften the refusal, i.e. in accordance with negative politeness strategy 6, the speaker may give unrequested information about the cause as an excuse.

Example (2)[2] illustrates the **Answer-ref-indicated** stimulus condition.

```
2.  Q: You're on that?
    R: No no no.
       Dave is.
```

Answer-ref-indicated: R suspects that Q is asking an indirect *wh*-question.

Again, techniques for interpreting indirect speech acts can be used to determine when this condition obtains. Justification for the stimulus condition can be made on grounds of cooperation, since providing the unrequested information saves Q the effort of making a *wh*-question as a follow-up, and on grounds of politeness, by positive politeness strategy 1.

In example (3)[3], which illustrates the **Clarify-extent-indicated** stimulus condition, R provides qualifying information to avoid an erroneous scalar implicature (Hirschberg 1985).

```
3.   Q: Have you gotten the letters yet?
     R: I've gotten the letter from X.
```

In other words, if R has gotten some but not all of the letters, a Yes would be untruthful. However, a No might lead Q to erroneously conclude that R has not gotten any of the letters. We characterize this stimulus condition in terms of Hirschberg's theory of scalar implicature.

Clarify-extent-indicated: R believes that there is a q such that q is the highest-ranked element of a set of salient possible truthful responses, and R suspects or believes that Q does not know that q holds.

Example (4)[4] illustrates the **Appeasement-indicated** stimulus condition.

```
4   Q: Did you wash the dishes?
    R: I brought you some flowers.
```

Appeasement-indicated: R suspects that the answer would be undesirable to Q but that some other information would be desirable to Q.

[2](Stenström 1984), example (102).

[3](Hirschberg 1985), example (59)

[4]Constructed example.

The satisfaction of this condition could be detected using heuristic rules and stereotypical and specific knowledge about Q's desires. Note that this stimulus condition can be justified by positive politeness strategies 1 and 6.

Providing Unrequested Information

Overview

This section describes the declarative knowledge and algorithm used in our system which enables the system to provide unrequested but relevant information in a response. (For details, see (Green 1994; Green & Carberry 1994).) The algorithm is to be invoked whenever it is the system's turn and the strongest discourse expectation (Carberry 90) is that the system will provide an answer to a Yes-No question, that is, to a request to be informed of the truth of some proposition p.

The inputs to generation include

- the semantic representation of p
- a set of discourse plan operators
- a set of coherence rules
- a set of stimulus condition rules, and
- the system's beliefs, including
  - beliefs which are presumed to be shared with Q (which we refer to as *shared* beliefs.)
  - beliefs which are not presumed to be shared with Q (which we refer to as *non-shared beliefs*.

The output of the generation algorithm is a discourse plan which can be realized by a tactical generation component (McKeown 1985). We assume that when answer generation begins, the system's *only* goal is to satisfy the discourse expectation, and that the system can truthfully provide exactly one of the following answers based upon its non-shared beliefs: *yes* (that p is true), *no* (that p is false), *partly yes* (that p is partly but not completely true), *maybe yes* (that the system suspects that p is true), or *maybe no* (that the system suspects that p is false).

Discourse Plan Operators

Discourse plan operators provided in the model encode generic programs for expressing what we refer to as a *full answer*. A full answer consists of the direct answer and any additional relevant information which the system decides to provide on its own initiative. For example, the discourse plan operators for constructing full Yes (*Answer-yes*) and full No (*Answer-no*) answers are shown in Figure 1.

To briefly describe our formalism, variables are denoted by symbols prefixed with '?'. The first line of a discourse plan operator specifies the type of discourse act, the participants, and a propositional variable. For example, *(Answer-yes s h p)* describes the act of speaker s giving a Yes answer to the request made by hearer h to be informed of speaker's evaluation of the truth of the questioned proposition, p. *Applicability conditions* (Carberry 90) are conditions which must hold for an operator to be selected but which the planner will not attempt to bring about. For example, the second applicability condition of *Answer-no* requires that the speaker believe that p is false. The primary goals of a discourse plan specify the intended effects of the plan and do not play a role in the first phase of response generation, which is the focus of this paper.

The *nucleus* and *satellites*[5] describe communicative acts to be performed to achieve the goals of the plan. *Inform* is a primitive act that can be realized directly. The non-primitive acts are defined by discourse plan operators themselves. Thus, a discourse plan may have a hierarchical structure. A full answer may contain zero, one or more instances of each type of satellite specified in the operator. The satellites describe which types of unrequested but relevant information may be included in the response provided that certain restrictions (every applicability condition and at least one stimulus condition) are satisfied.

For example, consider the *Use-elaboration* and *Use-obstacle* discourse plan operators, shown in Figure 2, describing possible satellites of *Answer-yes* and *Answer-no*, respectively. In general, each satellite operator in our model has applicability conditions analogous to those shown in Figure 2. The first applicability condition of the *Use-elaboration* operator requires that the speaker believe that the relational proposition (Mann & Thompson 1983) *(cr-elaboration q p)* holds for some propositions q and p instantiating the variables $?q$ and $?p$, respectively. Informally speaking, this is a requirement that the speaker believe that q is an elaboration of p, where elaboration is one of the ways in which coherence is maintained. The second applicability condition requires this same proposition to be *mutually plausible*, so that the hearer will be able to recognize the intended relational proposition. The manner in which mutual plausibility is determined is described in the section on Coherence Rules.

In our model, applicability conditions are necessary but not sufficient for deciding to include a satellite in the response. Each satellite operator includes *stimulus conditions* (as described earlier in the paper). In order for an instance of a satellite to be included, all of the applicability conditions and at least one of the stimulus conditions of the satellite operator must be true. A full answer might be too verbose if applicability conditions provided the only constraint on including satellites. Furthermore, use of a goal-driven approach to selecting satellites would not result in the selection of any satellites at all, since we assume that at the time when the system is asked a question, the only goal is to provide the requested information. Stimu-

[5]This terminology was borrowed from Rhetorical Structure Theory (Mann & Thompson 1987), which will be discussed briefly in the section on Coherence Rules.

```
(Answer-yes s h ?p):
  Applicability conditions:
    (discourse-expectation
            (informif s h ?p))
    (bel s ?p)
  Nucleus:
    (inform s h ?p)
  Satellites:
    (Use-condition s h ?p)
    (Use-elaboration s h ?p)
    (Use-cause s h ?p)
  Primary goals:
    (BMB h s ?p)
```

```
(Answer-no s h ?p):
  Applicability conditions:
    (discourse-expectation
            (informif s h ?p)
    (bel s (not ?p))
  Nucleus:
    (inform s h (not ?p))
  Satellites:
    (Use-otherwise s h (not ?p))
    (Use-obstacle s h (not ?p))
    (Use-contrast s h (not ?p))
  Primary goals:
    (BMB h s (not ?p))
```

Figure 1: Discourse Plan Operators for Yes and No Answer

```
(Use-elaboration s h ?p):
  Existential variable: ?q
  Applicability conditions:
    (bel s (cr-elaboration ?q ?p))
    (Plausible (cr-elaboration ?q ?p))
  Stimulus conditions:
    (answer-ref-indicated s h ?p ?q)
    (clarify-concept-indicated s h ?p ?q)
  Nucleus:
    (inform s h ?q)
  Satellites:
    (Use-cause s h ?q)
    (Use-elaboration s h ?q)
  Primary goals:
    (BMB h s (cr-elaboration ?q ?p))
```

```
(Use-obstacle s h ?p):
  Existential variable: ?q
  Applicability conditions:
    (bel s (cr-obstacle ?q ?p))
    (Plausible (cr-obstacle ?q ?p))
  Stimulus conditions:
    (explanation-indicated s h ?p ?q)
    (excuse-indicated s h ?p ?q)
  Nucleus:
    (inform s h ?q)
  Satellites:
    (Use-obstacle s h ?q)
    (Use-elaboration s h ?q)
  Primary goals:
    (BMB h s (cr-obstacle ?q ?p))
```

Figure 2: Two satellite discourse plan operators

lus conditions can be thought of as situational triggers which give rise to new speaker goals, the primary goals of the associated satellite operator, and which are the compiled result of "deeper" planning based upon intentional goals.[6] Note that the situation triggering a stimulus condition may be either the actual situation before the system has begun constructing a response, or the *anticipated* situation which would result from providing part of the planned response.

Coherence Rules

Coherence relations in discourse have been described by (Mann & Thompson 1983; Mann & Thompson 1987) and others. We claim that certain coherence relations can be used to characterize various types of satellites of full answers (Green 1994). Our model includes a set of *coherence rules*, encoded as Horn clauses, which provide sufficient conditions for the mutual plausibility of a relational proposition *(CR q p)*, where *CR* is a coherence relation and *q* and *p* are propositions. If the relational proposition is plausible to the system with respect to the beliefs which it presumes to be shared with the questioner, we assume that it would be plausible to the questioner too.

Stimulus condition rules

The stimulus conditions listed in satellite operators are evaluated using a set of stimulus condition rules, encoded as Horn clauses. The rules provide sufficient conditions for determining if the given stimulus condition holds in the current discourse context. Several of the rules were described informally above.

Algorithm

Content planning is performed by top-down expansion of an answer discourse plan operator. The process begins by instantiating each top-level answer discourse plan operator with the questioned proposition until one is found such that its applicability conditions hold. Then, this operator is expanded. A discourse plan operator is expanded by deciding which of its satellites to include in the full answer and expanding each of them (recursively).

The algorithm for expanding satellites is given in Table 1. Each stimulus condition of the satellite is considered in turn. If a conjunction of the applicability conditions and the stimulus condition holds for some instantiation of the variables of the satellite, then that instance of the satellite is included in the plan. Thus, the same type of satellite may be included more than once, although for different reasons, in the same response.

[6] As we discuss in (Green & Carberry in preparation), it is possible to translate the knowledge encoded as stimulus conditions in our model into an alternate formalism in which the intentional subgoals (Moore & Pollack 1992) of the discourse plan are represented explicitly in addition to the rhetorical subgoals provided in the current formalism.

Table 1: Algorithm for Expanding a Satellite

INPUT	sat-header:	satellite operator header, (Use-CR s h p)
OUTPUT	sat-plans:	list of instances of discourse plan operator with header identical to sat-header, or null

1. Set sat-op to operator, (Use-CR s h ?p), with ?p instantiated with p from sat-header.
2. For each stimulus condition SC of sat-op, create formula, PAT, which is conjunction of sat-op's applicability conditions and SC:
   - 2.a. If PAT is provable, then:
     - 2.a.1. Create instance, I, of sat-op using the assignment for the existential variable in PAT,
     - 2.a.2. Expand I's satellites, and
     - 2.a.3. Add I to sat-plans.
   - 2.b. If PAT is not provable for any stimulus condition, then sat-plans is null.

Example

To give a brief example of how unrequested information is selected, consider (5).[7] Although the full response shown in (5.a) - (5.e) would be selected during the first phase of response generation, the parts of the response appearing in square brackets would not be given explicitly by the system if (during the second phase of response generation) the system determines that they are inferrable from the rest of the response. That is, our complete model accounts for the generation of (5.d) - (5.e) alone as an indirect answer.

```
5.    Q:  Is Mark here [at the office]?
5.a   R:  [No.]
5.b       [He is at home.]
5.c       [He is caring for his daughter.]
5.d       His daughter has the measles.
5.e       But he is logged on.
```

First (5.a), which provides the requested information, is selected as a consequence of instantiating an *Answer-no* operator (given in Figure 1) in order to satisfy the current discourse expectation. This operator specifies that a coherent No answer may include additional information which contrasts with the information conveyed by (5.a), i.e., information resulting from expanding a *Use-contrast* satellite (given in Figure 2). In order for an instance of *Use-contrast* to be included in the discourse plan, the system must be able to instantiate the satellite operator's existential variable with a proposition *q* satisfying the applicability

[7] (5.d) was observed by one of the authors in a naturally occurring exchange, and (5.e) has been added for expository purposes.

conditions and at least one of the stimulus conditions. A search of the system's knowledge base of beliefs finds a proposition realizable as (5.b) for the stimulus condition of *Answer-ref-indicated*, and a proposition realizable as (5.e) for the *Substitute-indicated* stimulus condition. In other words, (5.b) is included because the system suspects that Q wants to know where Mark is, and (5.e) is included because the system suspects that the information expressed as (5.e) would help Q to overcome an obstacle. Thus, two *Use-Contrast* satellites are added to the top-level of the plan.

The *Use-Contrast* satellite whose nucleus is realized as (5.b) can be expanded further: (5.c) is provided as a *Use-Cause* satellite of it motivated by the stimulus condition *Explanation-indicated*, i.e., because the system suspects that Q will be surprised by the information given in (5.b). Similarly, a *Use-Cause* satellite is added to this *Use-Cause* satellite because the system sustpects that Q will be surprised by the information given in (5.c).

In summary, each of the unrequested pieces of information can be seen as resulting from the system's beliefs about Q and the effect that other parts of the planned response may have on Q. The following imaginary dialogue between Q and R illustrates the role of stimulus conditions in deciding what unrequested information to include.

```
6.     Q:  Is Mark here at the office?
      [Q:  Where is he? I need to contact him.]
6.a   [R:  No, he's not here.]
      [Q:  Well where is he?]
6.b   [R:  He is at home.]
      [Q:  Why?]
6.c   [R:  He is caring for his daughter.]
      [Q:  Why?]
6.d    R:  His daughter has the measles.
      [Q:  I need to contact him.]
6.e    R:  He is logged on.
```

In this example, the system has no other information on the subject and so will not provide any other information. An interesting area for future research is whether, given an agent's shared and nonshared beliefs, the applicability conditions and stimulus conditions sufficiently constrain the amount of unrequested information selected or whether other types of constraints should be incorporated (e.g., attentional, stylistic) into our model. Also, it remains to be seen whether different stimulus conditions should be given different priorities.

Relation to Previous Work

Some previous work in cooperative response generation describes discourse contexts in which certain types of unrequested information should be given in order to avoid misleading the questioner (Hirschberg 1985; Joshi, Webber, & Weischedel 1984) or to help the questioner overcome obstacles to his inferred plans (Allen & Perrault 1980). Contexts described in that work relevant to the generation and interpretation of answers to Yes-No questions have been represented in our model as stimulus conditions (e.g., *Clarify-extent-indicated*, *Answer-ref-indicated*). In addition to identifying other important stimulus conditions (e.g., *Explanation-indicated* and *Excuse-indicated*), our approach to providing unrequested information differs from past work in two ways. First, we provide a unified model for taking the initiative to provide this type of unrequested information. Second, past work was limited to generating single-sentence responses sensitive to the actual discourse context. Our model is capable of generating multiple-sentence responses containing unrequested information addressing the actual discourse context as well as anticipating the effect of parts of the planned response on the addressee.

Conclusions

We present an approach to taking the initiative to provide unrequested but relevant information in responses to Yes-No questions. In our model, stimulus conditions represent types of discourse contexts in which a speaker would be motivated to provide unrequested information. Stimulus conditions can be thought of as situational triggers which give rise to new speaker goals, and which are the compiled result of "deeper" planning based upon intentional goals. Stimulus conditions are used as part of a discourse-plan based approach to response generation which enables the generation of coherent multiple-sentence responses, including indirect answers.

References

James F. Allen and C. Raymond Perrault. Analyzing intention in utterances. *Artificial Intelligence*, 15:143–178, 1980.

Penelope Brown and Stephen Levinson. Universals in language usage: Politeness phenomena. In Esther N. Goody, editor, *Questions and politeness: Strategies in social interaction*, pages 56–289. Cambridge University Press, Cambridge, 1978.

Sandra Carberry. *Plan Recognition in Natural Language Dialogue.* MIT Press, Cambridge, Massachusetts, 1990.

Nancy Green and Sandra Carberry. Interpreting and generating indirect answers. In preparation.

N. Green and S. Carberry. Generating indirect answers to yes-no questions. In *Proceedings of the Seventh International Workshop on Natural Language Generation*, pages 189–198, Kennebunkport, Maine, 1994.

Nancy L. Green. *A Computational Model for Generating and Interpreting Indirect Answers.* PhD thesis, Department of Computer and Information Sciences, University of Delaware, 1994.

H. Paul Grice. Logic and conversation. In P. Cole and J. L. Morgan, editors, *Syntax and Semantics III: Speech Acts*, pages 41–58. Academic Press, New York, 1975.

Elizabeth Ann Hinkelman. *Linguistic and Pragmatic Constraints on Utterance Interpretation.* PhD thesis, Department of Computer Science, University of Rochester, 1989.

Julia Bell Hirschberg. *A Theory of Scalar Implicature.* PhD thesis, Department of Computer and Information Science, University of Pennsylvania, 1985.

Aravind Joshi, Bonnie Webber, and Ralph M. Weischedel. Preventing false inferences. In *Proceedings of Coling84*, pages 134–138, Stanford University, California, July 1984. Association for Computational Linguistics.

S. Levinson. *Pragmatics.* Cambridge University Press, Cambridge, 1983.

Kathleen R. McKeown. *Text Generation.* Cambridge University Press, 1985.

Johanna D. Moore and Martha E. Pollack. A problem for RST: The need for multi-level discourse analysis. *Computational Linguistics*, 18(4):537–544, December 1992.

W. C. Mann and S. A. Thompson. Relational propositions in discourse. Technical Report ISI/RR-83-115, Information Sciences Institute, University of Southern California, Marina del Rey, California, November 1983.

W. C. Mann and S. A. Thompson. Rhetorical structure theory: Toward a functional theory of text organization. *Text*, 8(3):167–182, 1987.

R. Perrault and J. Allen. A plan-based analysis of indirect speech acts. *American Journal of Computational Linguistics*, 6(3-4):167–182, 1980.

Anna-Brita Stenström. Questions and responses in English conversation. In Claes Schaar and Jan Svartvik, editors, *Lund Studies in English 68.* CWK Gleerup, Malmö, Sweden, 1984.