

SRI International



WHY ASK?

Technical Note 169

October 6, 1978

By: Jerry R. Hobbs, Computer Scientist
Jane J. Robinson, Senior Research Linguist
Artificial Intelligence Center

SRI projects 5844, 6891, and 7500

The work reported herein was supported by the National Science Foundation under Grants No. MCS 76-22004 and MCS 78-07121 and by the Advance Research Projects Agency under Contract N00039-78-C0060 with the Naval Electronic Systems Command.

ABSTRACT

In this paper, we address the problem, "What makes an answer appropriate?" We do so by investigating indirect answers to questions in task-oriented dialogues. Three cases are distinguished: (1) The response, though indirect, answers the question asked; (2) the response denies a presupposition of the question; and (3) the response answers to higher goals the questioner was trying to achieve. Detailed analysis shows the need for knowledge about the task, the role of the participants, and communication goals, in the construction of appropriate answers. We conclude with a preliminary formulation of the appropriateness of an answer in terms of the goals of the questioner and the knowledge of the respondent.

Is it possible, given a question in a conversation, to describe in a precise fashion, what would count as an answer? We are investigating this problem by analyzing a particularly difficult set of cases -- question-answer sequences where the answers are indirect. Our claim is that an appropriate answer must address the goals of the question, whether directly or indirectly. This accounts for our title, "Why Ask?"

We have distinguished three cases of indirect answers: First, there is the case where the answer is indirect but it does answer the question that was asked. In the second case, the answer denies the presupposition of the question. In the third, the answer addresses not the question itself but some higher goal the speaker was trying to achieve through asking it. Intuitively, we can see in all three cases that these indirect answers are appropriate, informative, and coherent. But what makes them so? In this paper we point out some of the kinds of knowledge that must be taken into account and illustrate the sort of analysis that needs to be applied to show the answers' appropriateness.

Our approach to the problem is to analyze utterances as forms of purposeful behavior -- behavior directed towards achieving goals through the use of language.* But this approach requires us to identify the purposes that lie behind someone's saying something. People are usually quite good at discerning these purposes, but how they do it is still something of a mystery in spite of the insights provided by theories of speech acts and conversation (Labov & Fanshel 1977, Searle 1969, Grice 1975). There are several complications. For one, goals are not fixed; they shift and sometimes conflict. For another, a remark may serve more than one purpose simultaneously, as when someone says "When do you plan

* Lehnert (1977) also investigates indirect answers to questions, but not in terms of the questioner's goals. Rather, she classifies the question into one of several categories and bases the response on this classification. Cohen (1978) and Allen and Perrault (1978) use the presumed goals of questioners to infer the omitted material in elliptical questions so as to be able to frame a direct answer, but they do not consider the problem of indirect answers. Harrah (1963) and Belnap and Steel (1976) present ways of formulating a logical representation of questions, when the goals or "interests" of the question are known.

to come home?" (Labov & Fanshel 1977) in order to obtain information and request the hearer to perform the action of coming home. Finally, the speaker's goals are generally inaccessible to the linguist. Introspection about what the goals might have been is highly speculative, and eliciting the speaker's report of his goals yields data with a rather dubious relation to the original utterance.

To control such factors, we have restricted our attention to dialogues between an Expert and an Apprentice engaged in repairing an appliance (Grosz 1977). To limit reliance on non-linguistic cooperation, the two participants were placed in different rooms, they communicated via teletype, and the Expert was given only a limited view of the task. This experimental setup overcomes some of the difficulties noted above by providing a 'real world' context with explicit, agreed-upon, well-structured goals, involving alternative states and sequences of actions that can be defined precisely. Our interpretations of the dialogues employ models of planning behavior that have been developed by workers in artificial intelligence (Sacerdoti 1977). Using these models, we can interpret precisely utterances involving modals like "can" and "should" and notions of obligation, possibility, and future states. In the task context, the assertion "You should take off the pump belt next" or the imperative "Take off the pump belt" are not moral value judgments or expressions of personal desire so much as they are expressions with propositional content that may be true or false or inappropriate at a particular time in a particular performance of the task.

In addition to the explicit goal of repairing the appliance, there are two implicit ones. One is implicit in the roles of the participants, the Expert who is to guide and teach, and the Apprentice, who is to learn the general principles of repair and maintenance of appliances through performing the task. The other implicit goal for both is to communicate information successfully in order to achieve the other two goals (cf. Robinson 1978). In order to achieve that goal, the Expert and Apprentice had to ask each other questions, the

Apprentice to find out how to proceed, the Expert to assess the Apprentice's progress. These questions and their responses have provided our material.

In the first of the three cases of indirect responses that we distinguish, the answer, although indirect, does answer the question that was asked. An example is:

- (1) E: Have you disconnected the air line?
A: I loosened it.

This is a cooperative conversation, so Grice's maxim of quantity (Grice 1975) -- "Make your contribution as informative as is required!" -- is applicable, and we are entitled to assume that the Apprentice did no more than he said. Now if "to disconnect" means "to cause not to be attached" and "to loosen" means "to cause not to be tightly attached", then although the air line is not attached tightly, it is still attached, and the literal answer is therefore "No". The actual response is indirect -- the literal answer must be inferred -- but the response answers the question. This is in contrast with examples that follow.

The second case involves questions that have no literal answers because the presuppositions of the questions are not true.*

- (2) A: Should I turn it [the faucet] off or leave it on?
E: It doesn't really matter.

Alternative questions presuppose that one of the alternatives is true. In (2), the Apprentice is hazarding a guess that one of two actions ought to be performed; the presupposition is "Either I should turn it off or I should leave it on." The task model allows us to give a precise interpretation to the modal "should": We can say that some action should be performed at a particular time if it leads to a state that is closer in the task model to the goal of the task, or if not doing the action leads to a state farther from the goal. But in the

* Kaplan (1977) has looked at the problem of generating responses that don't mislead when an existential presupposition of a question is false.

situation of (2), neither turning it off nor leaving it on leads to a better state, so the presupposition is false and the Expert denies it, setting the Apprentice straight as to what is relevant to the task goal.

There remain a large number of exchanges in the third class, in which the answer, though appropriate in context, does not answer the question posed. The three examples we present below all come from the stretch of continuous dialogue shown in Figure 1.

- A: What is the next thing to be attached?
E: Bolt the pump to the base plate.
. . . .
E: What tools are you using?
A: My fingers.
E: You will need to use two 1/2" wrenches, one on the top to hold the bolt and one underneath to tighten the nut.
. . . . [Long interval of silence]
(3) E: What are you doing?
(4) A: I would like to know if I can take off the back plate.
E: You shouldn't have to.
(5) E: Are you having trouble with the bolts?
A: Yes.
E: Use the ratchet wrench on the top and hold the nut stationary on the bottom with a box wrench.
(6) A: What's a ratchet wrench?
E: Show me the table. (TV camera provides vision)
The ratchet wrench is the object lying between the wheel puller and the box wrenches on the table.

Figure 1

First consider exchange (3):

- (3) E: What are you doing?
A: I would like to know if I can take off the back plate.

At any given moment, one is "doing" many things, so there are many ways one can answer the question, "What are you doing?" It is instructive to look at some of the possibilities. The Apprentice could have said

"I'm breathing,"

which would have been correct, but would have had nothing to do with the task. She could have said

"I'm fixing the air compressor,"

or

"I'm turning a wrench."

These would have been task-related, but they still would have been unsatisfactory.

The Expert's question is prompted by a passage of time with no report on the Apprentice's progress. It arises out of the Expert's need, as guide, to monitor the task. The Apprentice has to realize that the goal of the Expert's question is to inquire, in a more general way, "What success are you having in carrying out the current subtask?" A response such as

"I'm trying to tighten the back bolts,"

or

(7) "I'm having trouble with the bolts,"

would have been direct and to the point, and at first glance may even seem more appropriate than the answer actually given.

However, if we look back at the Apprentice's response in context, we can see that she is guessing what action would be appropriate. Because of the maxim of quantity, it is very common in seeking information or indicating a difficulty, for one to frame questions in a way that indicates how much has already been figured out, and this often includes guesses as well. Therefore, to answer as in (7) would actually have been less informative, because it would not have included the alternatives the Apprentice was considering. We can see this same principle at work in the Apprentice's question in (2), where the Apprentice guessed that it might matter whether the faucet was on or off.

The Apprentice's answer in (4) is itself an indirect question that the Expert answers indirectly:

- (4) A: I would like to know if I can take off the back plate.
E: You shouldn't have to.

The presupposition is "I can take off the back plate or I cannot take off the back plate," which is tautologically true. There is no problem of the presupposition's being false, and a direct answer, either "Yes" or "No", exists. There is a problem with the interpretation of the modal "can". Is the Apprentice asking for permission or is she asking about physical possibility? A plausible interpretation is that she is asking about possibility as an indirect way of asking for permission, since possibility is a precondition for felicitous granting of permission. However, it is a waste of time for the Expert to resolve the ambiguity. The Expert is entitled to assume the Apprentice's question is relevant to the task of bolting the pump to the base plate. Without resolving the ambiguity of "can" or replying directly, the Expert denies the necessity for the action, while leaving open the question of its possibility or permissibility. This informs the Apprentice that it is possible to get the job done without removing the back plate.

But negative answers are notoriously unsatisfactory. If the Expert has indirectly informed the Apprentice that she is on the wrong track, why doesn't he tell her what the right one is? The next exchange indicates that he is aware of her need for a more positive response. In order to formulate it, he must assume that her question is relevant to the task although the action she proposed is not. The Expert hazards a guess as to what gave rise to the request, and the guess is verified.

- (5) E: Are you having trouble with the bolts?
A: Yes.

The Expert's question in (5) is actually part of his response to the Apprentice's question in (4).

Our final example illustrates other considerations that must be brought to bear in the interpretation of an answer:

- E: Use the ratchet wrench on the top....
(6) A: What's a ratchet wrench?
E: The ratchet wrench is the object lying between the wheel puller and the box wrenches on the table.

Taken in isolation, the Apprentice's question is generic. It is a request for a definition, and the Expert could have answered directly by saying,

"A ratchet wrench is a wrench with a pawl, or hinged catch, that engages the sloping teeth of a gear, permitting motion in one direction only."

Instead he answers by giving the location of a specific ratchet wrench. Why does the answer satisfy the requirements of relevance and informativeness? We may conjecture two paths that could have led to this response -- the first a "task-based" analysis resting on the fact that the Expert has understood the problem that motivated the Apprentice's question, and the second a "text-based" analysis relying on a general conversational convention.

The first path is based on the fact that the Apprentice is in a situation where she must use the ratchet wrench. But to use something, if one does not already have it, one must get it. To get it, one must know or come to know its location, that is, one must find it. To find it, one must be able to recognize it. Since the Apprentice is unable to recognize it, she asks for its identifying characteristics -- "What is a ratchet wrench?"

One possible explanation of the Expert's response is that he understands the chain of preconditions that led to the Apprentice's question, and rather than answering at the end of the chain by giving the definition, he answers to the middle by giving the location, thus obviating the search and recognition problems. In a sense, we can say the Expert gave the most informative answer under the circumstances, because he addressed the chain at a point as close as possible to the ultimate goal of "using", in spite of the fact that the question in isolation would have been answered more directly by the definition. Had he been present in the room, he could have simply handed it to her,

thereby causing her to have it without satisfying any of the other preconditions for using it.

The second, text-based, path to an interpretation of the exchange is via a conversational convention. It is one of Grice's conversational maxims that participants in a conversation are to "Be relevant!" Although Grice himself points out the difficulties in defining relevance, or coherence, in conversation, there are some ways of continuing a conversation coherently that are so common that people are "primed" to recognize them. One such continuation serves the goal of successful communication. Let us call it the Request for Repair of a Definite Reference (cf. Jefferson 1972, Hobbs 1978).

Ideally, a speaker will choose his descriptions of the entities involved in his utterance in such a way that the listener will be able to identify the entity referred to. That is, he will use properties the listener knows about and words the listener is familiar with. Clark and Marshall (1978) point out further that to be felicitous, a definite reference must use only properties in the speaker's and hearer's domain of mutual knowledge; that is, the speaker must believe the properties are true of the entity, and must believe the hearer believes it, and must believe the hearer believes the speaker believes it, and so on, ad infinitum. In actual performance, however, speakers tend to be quite casual in choosing their descriptions, for since there is interaction, there is always an opportunity to repair. That is, the speaker will frequently describe the entity in the way that is most natural for him, and will take pains with the description only if it turns out that the initial attempt is not understood. The tendency to "under-describe" is further heightened by the fact that to "over-describe" is frequently taken as a challenge to the listener's status as an intelligent adult. All of this implies that a speaker must be ready to interpret a response to something he says as a Request for Repair of Definite Reference.

In (6), the Apprentice is asking for further identifying properties of the entire class of ratchet wrenches. But an answer to this question would give her further identifying properties of the particular wrench

referred to by the Expert. Since the question would not otherwise be relevant, the Expert may assume it to be a request for repair of his definite reference to the ratchet wrench. In general, one may repair a definite reference by giving whatever further identifying characteristics one thinks best, regardless of what a direct answer to the request would be. Thus, instead of answering the question that is asked, the Expert decides to give another uniquely identifying property -- its location -- and informs the Apprentice of that.

The second, text-based analysis is like the first, task-based analysis in that for both, the Expert must get behind the Apprentice's question to see the problem she is facing. However, here in the text-based analysis, the problem is not the task problem of using the ratchet wrench, but the communication problem of understanding the referring expression "the ratchet wrench".

As a result of these investigations, we propose the following preliminary answer to the question "What makes an answer appropriate?" Utterances that function as questions typically request information that the questioner believes will help him achieve some goal; an answer is appropriate if it provides some information that allows him to achieve that goal, although perhaps not the information he asked for. If the respondent does not have information sufficient for achieving the goals, his answer will be most appropriate if it includes the information that contributes most toward the goals. This formulation will doubtless have to be refined to take care of such instances as when the respondent disapproves of the questioner's goals, has good reason to want the questioner to figure out a solution to his problem from partial information, and so forth.

The formulation also suffers from the problems of indeterminacy of the participants' beliefs and goals and from the vagueness of the phrases "most appropriate" and "contribute most". In general, it is hard to avoid these problems. It is for this reason that task-oriented dialogues have proven especially useful as a site for a preliminary

study of question-answer sequences, as indeed for discourse analysis in general; for here the beliefs and goals of the participants can be stated objectively and with some precision.

REFERENCES

- Allen, J. & C. Perrault 1978. Participating in dialogues: Understanding via plan deduction. Proceedings, Second National Conference, Canadian Society for Computational Studies of Intelligence, Toronto.
- Belnap, N. & T. Steel 1976. The logic of questions and answers. Yale University Press, New Haven.
- Clark, H. & C. Marshall 1978. Definite reference and mutual knowledge. Paper presented at the Sloan Workshop on Computational Aspects of Linguistic Structure and Discourse Setting. University of Pennsylvania, May 1978.
- Cohen, P. 1978. On knowing what to say: planning speech acts. Technical Report No. 118, Department of Computer Science, University of Toronto. January 1978.
- Grice, H. 1975. Logic and conversation. In D. Davidson and G. Harman (Eds.), The Logic of Grammar. Dickenson. Encino, California.
- Grosz, B. 1977. The representation and use of focus in dialogue understanding. Stanford Research Institute Technical Note 151, Stanford Research Institute, Menlo Park, California.
- Harrah, D. 1963. Communication: A logical model. M.I.T. Press, Cambridge, Mass.
- Hobbs, J. 1978. Coherence and coreference. SRI Technical Note 168, SRI International, Menlo Park, California.
- Jefferson, G. 1972. Side sequences. In D.N. Sudnow (Ed.) Studies in social interaction, 294-338. Free Press, New York.
- Kaplan, S.J. 1977. Cooperative responses from a natural language data base query system. Department of Computer and Information Science, University of Pennsylvania, Philadelphia, Pennsylvania.
- Labov, W. & D. Fanshel 1977. Therapeutic discourse. Academic Press, New York.
- Lehnert, W. 1977. A conceptual theory of question answering. Proceedings, IJCAI-77, Cambridge, Mass.

Robinson, J. 1978. Purposeful questions and pointed answers. Presented at 7th International Conference on Computational Linguistics, Bergen, Norway.

Sacerdoti, E. 1977, A structure for plans and behavior. Elsevier, New York.

Searle, J. 1969. Speech acts: An essay in the philosophy of language. Cambridge University Press, Cambridge, England.