# USER PARTICIPATION IN THE REASONING PROCESSES OF EXPERT SYSTEMS

Martha E. Pollack, Julia Hirschberg and Bonnie Webber

Department of Computer and Information Science
The Moore School
University of Pennsylvania, Philadelphia, Pa. 19104

We argue that expert systems, if they are to satisfy the legitimate needs of their users, must include dialogue capabilities as sophisticated as those proposed in current Natural Language research. In particular they must allow the user to direct the flow of the dialogue and must take into account the user's goals and expectations both in analyzing the user's statements and in providing appropriate responses. Our claims are corroborated by an analysis of transcripts of a "naturally occurring" expert system, a radio talk show in which callers ask an expert for financial advice. We present data demonstrating that user-expert dialogues are best viewed as a negotiation process, and we describe the exchanges that compose the dialogue in terms of the motivations, goals, strategies, and moves of the participants.*

To date, user participation in the reasoning processes of expert systems has been largely limited to probing expert reasoning or adding limited information. The user may only ask why the system requested more information and how it arrived at its advice. Research into extending the capabilities of expert systems [1,2,3] has so far failed to recognize the need to permit the full range of interactions possible when humans engage in the normal giving and getting of advice.

In order to delineate patterns of user-expert interaction in the reasoning process, we taped one week, February 1-5, 1982, of the radio talk show "Harry Gross: Speaking of Your Money" from station WCAU in Philadelphia. In this program listeners call in, introduce themselves and ask for financial advice, which the expert then attempts to provide. The ensuing discourse is a "naturally occurring" dialogue between a human user and a human expert. While the protocols we obtained should obviously not be seen as the ideal toward which all expert systems should aspire, despite the oral nature of this naturally occurring expert system, despite the entertainment element involved, and despite its domain specificity and the type of user toward which it is aimed, we believe that the major requirements of truly flexible and cooperative expert systems can be determined, in a principled way, from an explication of the relationship between human user input and human expert response.

From the radio show we collected twelve and one-half hours of user-expert interaction, involving 120 callers. An examination of these protocols reveals a regular pattern of interaction, which we describe as negotiation, the process whereby people arrive at a conclusion by means of a discussion. Rarely does a caller simply state a problem and passively listen to the expert's response. Rather the caller actively participates in the definition and resolution of the problem. Caller and expert must often negotiate to determine the statement of a problem the expert can solve and the statement of a solution the expert can support and the caller accept -- and, ideally, understand. They may also need to negotiate a common understanding of terminology, a common set of world or domain beliefs, or an acceptable justification for the solution.

Given this view of the interaction between user and expert as a negotiation process, one can characterize the particulars of each verbal exchange in terms of the participants' motivations, goals, strategies and moves . In order to generate a response appropriate to the user's participation, the expert must first recognize the user's motivation(s), goal(s) and strategy(s) through an examination of his/her move in the context of the discourse history and the user model, and then, based on that analysis and his own motivation(s), determine his own goal(s), strategy(s) and move.

When the user decides to participate s/he has some motive for doing so and some goal s/he hopes to achieve. Motivation answers the question "Why does the user decide to participate?", while goal answers "What does the user hope to achieve through participation?". The user then adopts some strategy by which s/he attempts to achieve the goal(s) s/he has set; so, strategy answers the question "How does the user attempt to achieve the goal?". This strategy is realized linguistically in an utterance, the user's move. For example, a user may be motivated by his/her surprise at an expert's advice: s/he may have thought of and rejected an answer because s/he believes it violates some specific constraint s/he wants met. His/her goal is then to gain assurance that the the expert's response meets that constraint. S/he might try to achieve this goal

*A longer version of this paper appears as University of Pennsylvania Technical Report #TR-CIS-82-9.

by highlighting the constraint; this she might do with the question "But that would be illegal (or immoral, or expensive, or fattening), wouldn't it?".

The expert, upon hearing the user's utterance, must recognize the user's intent in participating. Determining a user's strategy, goal and motivation is an important aspect of natural language understanding research; recent work on recognizing speaker's intentions [4,5] can help to explain this recognition process. However, we believe that in interactions with an expert system, the range of user motivations, goals and strategies is more constrained than in unrestricted discourse, thus facilitating the implementation of an expert system component to recognize intentions.

After recognizing the user's intentions, the expert must choose a suitable response and realize it linguistically. The latter problem has been studied in language generation research [6,7,8]. The former, deciding what response to provide to what form of participation is a goal of our further research.

The following exchange between a caller and the radio expert illustrates the negotiation process described above:

1. L: "Hi, Harry, this is L."
2. Expert: "Welcome, L."
3. L: "Oh, hi, I have a question about charitable deductions."
4. Expert: "Sure."
5. L: "We give, uh, let me see, my husband's income this past year is roughly $24,000, not quite, and we gave maybe $1400 to our church and other religious organizations that we have receipts for."
6. Expert: "Right."
7. L: "Now we also give to Sunday School, to the benevolent offering, to this and that that just is out of pocket money that we don't have receipts for. Given the amount of money that we give that we do have receipts for, what is the amount that we would like be able to claim on our income tax that IRS won't call us in?"
8. Expert: "Oh, you're asking for a figure that nobody can really give you. How much was that, do you know?"
9. L: "Well in the past we've declared around one hundred and they've never called us in, but we give almost double that."
10. Expert: "Well if you give that much, and you have some evidence that you give that much, that's fine. If you put $2 a week, for example, in an envelope, that comes to $100; if you put $3, you've got $150, and if that's your regular contribution, deduct it."
11. L: "But (pause) the thing is, like Sunday School offering they don't keep track of who gives what, and they have a benevolent fund when we have communion, you put $5, you don't keep track of who gives what, that's what I'm talking about."
12. Expert: "What I would suggest that you do, is hereafter keep a record of what you have given. All you can do now is estimate that amount to the best of your ability."
13. L: "Now my records would be okay now I mean (pause)"
14. Expert: "Absolutely, your records will stand up too."
15. L: "Oh, I didn't know that!"
16. Expert: "Oh sure. What about the guy who goes out and spends money on a taxi cab and leaves the guy a tip? That tip is deductible and you don't get a receipt for a tip. You might even not get a receipt from the taxi cab."
17. L: "Oh, I didn't know that. I thought it was just what we had a receipt from the church for that (pause) "
18. Expert: "Oh, no."
19. Expert: "Plus the fact, the mere fact that the IRS calls you in and even desires to disallow part of your deduction, you're not required to agree with them."
20. L: "Right (pause) "
21. Expert: "You stand up for your rights. Whatever you give to charity, you claim."
22. L (laughing): "I don't want the hassle, of an, of an (pause) "
23. Expert: "Well, now you've got another point, and I think at times the service counts on the fact that people don't want the hassle. And maybe we as Americans have to stand up a little bit more and claim what's due us."
24. L: "Okay, I'll keep my own record of it from here on in."
25. Expert: "That's right, and estimate the best that you can based on what you did in the past, and your record of what you do in 1982 if you're ever called in will assist you in that. Okay?"
26. L: "Okay, now in the past we've declared $100 of un -- "
27. Expert: "But if you feel it's closer to $200, use the $200."
28. L: "Now could I jump from $100 to $200 in one year?"
29. Expert: "Absolutely!"
30. L: "I could?"
31. Expert: "Absolutely!"
32. L: "Okay?"
33. Expert: "All right?"
34. L: "Okay, thank you."

In (1-2) L and the expert initiate the dialogue by identifying themselves. This mutual introduction is important in beginning the process by which the expert develops a model of the caller -- and the caller of the expert. In the radio show it often establishes whether or not the caller has previously used the system, i.e., is already known to the expert. It also establishes initial rapport between caller and expert. The exchange in (3-4) shows L limiting the domain of her query before even posing it. The expert recognizes in L's initial participation in the reasoning process an implicit desire for reassurance that her question will be appropriate, and he provides such reassurance in (4). In (5) L continues to participate in the reasoning process by offering information she thinks the expert will need in order to answer her query -- still before the query is made: she believes her husband's

income and the amount of charitable donations for which they have receipts will be important factors in the expert's calculations. Her information adds to the expert's model of his caller, by providing family income level and by suggesting that L is not employed outside the home.

In (7) L finally poses her first direct query. She and her husband make charitable donations for which they do not have receipts and she wants to know, given the amount of donations for which they do have receipts, how much of the unreceipted donations they can deduct from their income tax. Note that in making this query, the type of query which existing systems might be expected to handle, L indicates two important points to the human expert which existing expert systems would not be capable of recognizing. First, she believes that the amount she can deduct for donations for which she has no receipts is a function of the amount of charitable donations for which she does have receipts ("Given the amount..."). Second, she imposes a user-specific constraint upon the expert: she does not want to be called in by the IRS to explain her deductions. This constraint is added to the expert's user model but not addressed directly until (19).

In response to (7), the expert first informs L that her initial query is unanswerable. However he goes beyond the simple rejection of her query to seek information which will help him answer what he perceives to be her real question: how can she deduct charitable donations for which she has no receipts. He must ascertain L's motivation behind a question which, as the expert says, no one can answer. Attempting to satisfy L's implicit goal, the expert begins to elicit information L had not realized was pertinent. First he asks for the amount of donations for which she has no receipts.(8). In providing this information in (9) L continues to participate in the reasoning process by indirectly reiterating her initial constraint: not being called in by the IRS. In (10) the expert gives a direct answer to what is in his view her true query: she can deduct the full amount she has donated. He also provides spontaneous justification for his response using a hypothetical suggested by his user model -- L is clearly a church-goer -- by noting that a regular, periodic contribution of a reasonable amount will be acceptable to the IRS even without written receipts.

In (11) L shows her dissatisfaction with his response, suggesting that the expert has not realized that the contributions she is worried about were not regular like the ones he noted in his hypothetical. She signals her dissatisfaction with the words "But (pause) the thing is..." and "...that's what I'm talking about." The expert's prior justification has led her to worry that he has failed to realize the facts in the case. She wants assurance that he has used all these facts in determining his advice, so she reiterates them. In (12) the expert recognizes her concern and provides a more suitable response: he accepts the information that her contributions are irregular and tells her she has no alternative but to estimate their amount. However, he then offers further advice for which L has not specifically asked: hereafter she should keep her own record of such donations to prevent a recurrence of her dilemma.

In (13-18) L and the expert negotiate further over the nature of IRS-acceptable records: note L's skepticism (13,15) which the expert recognizes as a request for additional reassurance. He provides this first by simple affirmation (14), next (16) by a hypothetical example and finally (18) by simple rejection of her preconception of (17).

In (19-23) the expert finally addresses L's previously expressed constraint of not being called in by the IRS. He tries to convince L that her constraint may be invalid by informing her that she need not accept an IRS challenge of her deductions (19). In response, L indicates through a pause after ostensible agreement (20) that she is unhappy with this argument, so the expert provides additional justification in (21): L has certain rights which she should exercise. However, in (22) L explains the purpose of her original constraint: she does not want the "hassle" of being called in by the IRS. The expert accepts her motivation as valid (23), but further justifies his advice by claiming that following it would benefit not just L but all who deal with the IRS.

In (24-31) L and the expert finish the discussion, reaching some agreement over how much of the expert's advice L intends to accept. In (24) she indicates her readiness to accept part of the expert's advice: that she keep her own records from now on. Note that the advice accepted is significantly different from the advice requested in her original query. The expert reiterates in (25) his further advice that she estimate for this year's income tax and adds the unsolicited information that her 1982 records will assist her if the IRS questions deductions in the future. In (26,28) L finally accepts the notion of increasing her estimate, but asks the expert for assurance that such a course of action will not cause problems. In (29-34) expert and user play out the final stage or end game of the negotiation process, one of expert reaffirmation, user desire for final reassurance, and expert concern that the user is truly satisfied with the advice given.

Clearly, existing expert systems cannot support such flexible exchanges. To do so requires a system capable of (i) detecting the user's explicit, implicit or indirectly stated goals, desires and expectations; (ii) performing reasoning consistent with them; and (iii) providing a satisfactory response. Most indirect user queries cannot be recognized as such by current systems. For example, the technique of reiterating a fact used by L above (11) is a

typical way for a human to signal to an expert a desire that he demonstrate his use of that fact; current systems require explicit user query. Similarly, reasoning consistent with indirectly stated goals and desires is impossible in current systems. Neither will current systems spontaneously answer questions not explicitly asked of them. When the expert realizes that L's concern is her lack of written records, he not only explains a viable alternative for the present, but he also provides additional information which will aid her in the preparation of future returns. Providing alternate forms of justification in response to user dissatisfaction is also unsupported by existing systems. In L's case, the expert's initial justification of the viability of personal records is unsatisfactory, while his subsequent statement that there is no alternative and, finally, his hypothetical example, eventually satisfy her.

The behavior exhibited by L is typical of the behavior we have observed in our protocols: people are often indirect in the their dealings with a human expert. If artificial expert systems are to become widely used, their designers must acknowledge the fact that people actively participate in the definition and resolution of their own problems. While people are not always able precisely to define the problems which cause them to seek information or advice, they do often have preconceptions about what a solution to their problems involves or what constraints it must satisfy. The negotiation model we have described provides a basis for characterizing what a user may be trying to do at any point in his/her interaction with an expert and identifies the options that are, in turn, available to the expert.

## ACKNOWLEDGEMENTS

We would like to thank Barbara Grosz and Aravind K. Joshi for their most helpful comments on this work.

## REFERENCES

[1] Swartout, W. R. "Producing Explanations and Justifications of Expert Consulting Programs," M.I.T. Technical Report, LCS TR-251.

[2] Clancey, W. J. "The Epistemology of a Rule-Based System: A Framework for Explanation," Stanford Technical Report, STAN-CS-81-896.

[3] Hayes, P., and Reddy, R. "An Anatomy of Graceful Interaction in Spoken and Written Man-Machine Communication," Carnegie-Mellon Technical Report CMU-CS-79-144.

[4] Cohen, P. R., Perrault, C. R. and Allen, J. F. "Beyond Question-Answering," Bolt Beranek and Newman Report No. 4644.

[5] Brown, G. P. "Characterizing Indirect Speech Acts," American Journal of Computational Linguistics, 6:3-4.

[6] McKeown, K. R. "Generating Natural Language Text in Response to Questions about Database Structure," Ph.D. thesis, Univ. of Penn., 1982.

[7] McDonald, D. D. "Natural Language Production as a Process of Decision Making under Constraint," Ph.D. thesis, M.I.T., 1980.

[8] Appelt, D. "Planning Natural-Language Utterances to Satisfy Multiple Goals," SRI Technical Note 259, 1982.