A SHORT NOTE ON OPPORTUNISTIC PLANNING AND MEMORY IN ARGUMENTS

Lawrence Birnbaum

Yale University
Department of Computer Science
New Haven, Connecticut

...chance favors only the mind that is prepared, - Pasteur

Introduction

Engaging in an argument is a complex task of natural language processing that involves understanding an opponent's utterances, discovering what his "point" is, determining whether his claims are believable, and fashioning a coherent rebuttal. Accomplishing these tasks requires the coordination of many different abilities and many different kinds of knowledge. Because arguing, and conversation generally, involve real-time interaction with another agent, this coordination must be even more flexible than is required for other natural language processing tasks. An arguer must have some expectations about what his opponent might say, but must also be able to respond to the unexpected. He must have some idea of the claims he wants to make, and plans for putting them forward, but his opponent may confound these plans. Or, more positively, his opponent may say something that offers an unforeseen opportunity to make a point. Arguing thus exemplifies the need for a flexible mix of top-down and bottom-up processing in both language understanding and production.

This paper is concerned with the roles of memory processing and planning in the processes of understanding and generating utterances in an argument or conversation. In particular, I will show that the memory and inferential processing necessary in order to understand another person's utterances can and should perform much of the work required to generate a response, work that most previous theories of conversation would delegate to explicit, goal-directed planning. The consequences of this, both for memory processing and for planning, will be briefly described and analysed.

The problem of choice

One of the most interesting, difficult, and yet frequently neglected questions that arises in analyzing a conversation or an argument is determining why a participant responds to a given utterance as he does, rather than in other equally plausible ways. In a *descriptive* theory of arguments, this problem need not be directly addressed, because the goals of such a theory would be satisfied if it were able to properly delineate the range of possible responses to an utterance in an argument. Such approaches typically characterise a response as the result of a series of hierarchically arranged choices among different response types (see, e.g., Reichman, 1081). For example, at the top level, the choices might be as follows:

- Attack one of the opponent's claims.
- Re-support one of your own claims.
- Change the subject.

In a *computational* theory, however, such a list can at most be only a first step towards solving the problem. Even if we characterise the conversationalist's task as involving a choice from among some such set of alternatives, a computational theory cannot sidestep the question, *how i$ the choice to be made?* The problem stems from the difficulty of choosing among alternatives in the absence of any reason to believe that the choices will actually lead to a good response. Too many decisions have to be made in the absence of any knowledge about what, ultimately, will be needed in the way of facts and reasoning about the topic domain to carry out a response based on those decisions. Such approaches must therefore rely heavily on back-up in order to produce useful responses. If back-up is to be avoided, then decision* about what to say should be based, as much as possible, on vhat is known about the subject under discussion.

The role of memory processing

It is by now a truism that memory and inference are central elements in natural language processing. In light of this central role, it seems likely that much of the explanation for conversational behavior arises from the contents and function of such a memory (Schank, 1077). In particular, memory plays a key role in explaining why a participant in a conversation or argument responds as he does, rather than in some other way. This claim is based on the observation that *a good response to an utterance in a conversation can often be discovered as a side-effect of the memory and inferential processing that is required simply in order to understand that utterance* (Birnbaum, Flowers, and McGuire, 1080).

For example, suppose that you were arguing with someone about the Viet Nam War, and he claimed that the Communists were responsible, because they refused to participate in the U.N.-sponsored elections intended to settle the political future of Viet Nam following the French pull-out. This claim happens to be false: It was the regime in the South, which, at the urging of the U.S., refused to participate in the elections. If you knew that fact, then you would undoubtedly notice that your opponent's claim was incorrect, and you would almost certainly say so in the context of an argument. You would probably also go on to point out how the facts in this case reflected on U.S. responsibility for the war.

However, you would probably also notice that the claim is

incorrect if you encountered it in a context other than an argument - for example, if you happened to read it in an account of a Presidential news conference. It seems obvious that if someone tells you something that you believe not to be the case, and you understand what he is saying, then you will notice the contradiction with your beliefs, regardless of whether or not you are engaged in an argument at the time. That is, the memory processing that uncovers such responses does not seem particularly argument-driven (although the decision as to whether or not to use them undoubtedly is). It is, more or less, the kind of processing that would be necessary to understand and assimilate the input regardless of whether the utterance occurred in an argument or in some other context. Memory processing must, as a matter of course, notice contradictions or inconsistencies between an input and the relevant beliefs of the understander. The utility of such processing in argumentation is obvious: If a contradiction is noticed in the course of understanding an opponent's utterance, then that contradiction is a good candidate to form the basis of a rebuttal. More generally, an increased reliance on memory allows the content of the discourse, and the speaker's knowledge, to play more active roles in the formation of response.

The role of opportunistic planning

Just as important in determining a speaker's response as his knowledge or the content of the discourse are his *goals* in the conversation or argument, and the plans by which those goals can be achieved. Indeed, much recent research on conversation has been based on the idea that conversation and other forms of discourse are planned behavior, in the same way as most other intelligent action is (see, e.g., Levin and Moore, 1077; Deese, 1078; Hobbs, 1070; Levy, 1070; Allen and Perrault, 1080). Most prior investigations of conversational behavior have employed some notion of planning borrowed from the problem-solving literature (see, e.g., Sacerdoti, 1077). All of these planning models are top-down in the rather straightforward sense that they all start with an explicit goal, and, with varying degrees of sophistication, attempt to devise a plan, or sequence of actions, that will satisfy that goal.

This top-down approach to planning can work in situations which are more or less under the control of the planner. Thus, it has proven useful in generating or understanding single utterances in a cooperative situation, as exemplified by the work on speech acts (Cohen and Perrault, 1070), or in conversations about tasks which themselves have a planned, hierarchical structure. But conversations and arguments do not, in general, meet those requirements. The actions of one's conversational partner, who in the case of an argument is assuredly not disposed to cooperate, can be expected to interfere with any top-down plan spanning several exchanges. Utterances in a conversation must not only further the speaker's own goals; they must also relate to what his partner (or adversary) has just said. Thus, unless a speaker can predict, rather specifically, how his adversary will respond, his utterances cannot be completely planned in advance.

The conclusion to be drawn, then, is that participating in a conversation or argument requires *opportunistic processing,* that is, both the ability to formulate plans, and the ability to recognise and pursue opportunities, in order to satisfy conversational goals (McGuire, Birnbaum, and Flowers, 1081). Obviously, the need for this kind of flexibility is not limited to conversations: It is a key factor in all kinds of intelligent behavior. To take an example from Meehan's (1070) TALE-SPIN domain, suppose that Joe Bear is hungry, and decides to ask Wilms Canary where some honey is. She tells Joe that she will answer him, if he brings her a worm. In the course of looking for a worm, Joe stumbles across some honey, or perhaps some fruit. If Joe does not eat the honey or the fruit, but rather continues searching for a worm to give Wilma, we would say that his behavior was not very intelligent. But without the ability to notice opportunities in the world that can be used to satisfy goals other than the one which is immediately governing his current behavior, that is exactly what would happen.

Recent research on planning and problem-solving has begun to address this point. Hayes-Roth and Hayes-Roth (1070) have proposed a model of opportunistic planning (the term is due to them) in which the planner's decisions in formulating a plan are not strictly hierarchical. Rather, decisions and observations at a given level of abstraction in the plan can influence not only the more specific levels that it dominates, but can also suggest opportunities to more abstract levels. Thus, for example, upon planning to go to the store to buy milk, the planner will notice that another of its currently active goals is to buy eggs. It will then plan to buy both milk and eggs at the store. (Hobbs, 1070, has pointed out the influence of something akin to this form of opportunism in conversational behavior.)

However, their model only exploits opportunities that arise while planning, not while executing plans. Thus, new goals cannot be formed as a result of noticing opportunities to achieve them. The difference between these two varieties of opportunism can be illustrated with another story from the shopping domain. Suppose a planner decides to go to the store to buy milk. In the course of executing that plan, while at the store, he notices that eggs are on sale. Realising that he will need eggs in the near future, he checks to see whether he has sufficient funds, and if he does, he buys some.

Wilensky's (1083) theory of *meta-planning* comes closer to satisfying the requirements for full opportunistic behavior. The most salient feature of the theory is its emphasis on goal *detection* as a central issue in planning and problem solving. That is, a planner must be able to figure out what its goals should be in a given situation. In Wilensky's model, the chief application of this ability is in noticing some interaction between several active goals (for example, that two goals are in conflict), and as a result formulating a new goal to deal with that interaction. However, the ability to determine what goals are relevant is crucial to opportunism as well. In order to exploit an opportunity, one must first recognise that it *is* an opportunity; and to do that, one must realise that the opportunity serves some goal that might be worth achieving. A goal-detection mechanism must therefore be able to suggest relevant goals upon the detection not only of problems, but also of fortuitous opportunities.

After an opportunity is noticed, the next step is clearly to decide whether or not to pursue the opportunity. In actuality, this must be a decision as to whether or not to pursue the goal that the opportunity can further. This in turn depends on what other goals the planner has active, including both previously active goals and those presented by other new opportunities, and of course, the time and resources that will be necessary in order to achieve each of these goals. In other

words, much of the effort required for opportunistic planning must be devoted to reasoning about goals: their desirability, their cost, and their interactions.

The application of these ideas to planning in arguments requires identifying argument goals, and strategies that choose among several possible goals. A very simple example of such a strategy might be to always exploit an opportunity to attack an erroneous factual claim of your opponent's. This strategy would be cheap to use, because most such opportunities would arise during the attempt to understand the erroneous claim. Thus, the work of determining that the claim was false, and why, would already have been accomplished. Because it is cheap, it would easily satisfy the general imperative of responding to your opponent's last utterance in a relevant way. If effectively pursued, the strategy would satisfy the goal of casting doubt on your opponent's credibility, and adding to your own. In this respect the strategy is quite aggressive. But in another sense, the strategy is rather passive: it would allow your opponent to set the agenda by making provocative claims. Thus, the unconditional use of this rule would reflect a decision, whether explicit or not, to set a higher priority on the goal of attacking the opponent than on the goal of controlling the topic of discussion.

Conclusions

In this paper, we have seen that conversational behavior, indeed intelligent behavior in general, requires the ability to seize opportunities to satisfy goals. In order to seize an opportunity, a planner must first be able to recognize one. That is, it must be able to recognize when a situation seems to facilitate the achievement of a worthwhile goal. Such an *opportunistic planner* must set its goals not only as a result of planning to achieve previously active, higher-level goals, but also by assessing the goals that its current situation presents opportunities to pursue.

The ability to recognize when a situation offers the fortuitous opportunity to achieve some goal puts a heavy burden on the perceptual, inferential, and memory capabilities a planner uses to understand and assess the situation in which it finds itself. To some extent, this burden can be eased if the planner has a general characterization of the sorts of opportunities that might arise in a given situation. For example, it seems reasonable that, when going to a grocery store, a planner should expect that some items may be on sale, although exactly which items will be unknown. However, in order to notice truly unexpected opportunities, a planner must be able to infer new goals from features of the situation not necessarily related to its currently active goals. This entails a fairly substantial capability for "bottom-up" inference.

The application of these ideas to conversational behavior offers the promise of a model that avoids many of the problems inherent in previous approaches. An opportunistic conversationalist would be able to set conversational goals in part on the basis of what its inferential processing uncovered in the course of understanding the content of the discourse. In particular, this would lead to a less top-down approach to the formation of responses, since the discovery of a potential response would determine which conversational goal to pursue as often as, if not more often than, the reverse. It is also worth noting here that the situations which provide these opportunities for response are themselves the result of the inferential memory processing that constitutes understanding. Opportunities exist not just in the world but in our thoughts.

Acknowledgments: I thank Mark Burstein, Gregg Collins, Natalie Dehn, Margot Flowers, Drew McDermott, Rod McGuire, and Roger Schank for useful discussions and for comments on earlier drafts of this paper. This work was supported in part by the Defense Advanced Research Projects Agency, monitored by the Office of Naval Research under contract N00014-85-K-0108.

References

Allen, J., and Perrault, C. 1080. Analyzing intention in utterances. *Artificial Intelligence,* vol. 15, pp. 143-178.

Birnbaum, L., Flowers, M., and McGuire, R. 1080. Towards an AI model of argumentation. *Proceedings of the First AAAI Conference,* Stanford, CA, pp. 313-315.

Cohen, P., and Perrault, C. 1070. Elements of a plan-based theory of speech acts. *Cognitive Science,* vol. 3, pp. 177-212.

Deese, J. 1078. Thought into speech. *American Scientist,* vol. 66, pp. 314-321.

Hayes-Roth, B., and Hayes-Roth, F. 1070. A cognitive model of planning. *Cognitive Science,* vol. 3, pp. 275-310.

Hobbs, J. 1070. Conversation as planned behavior. *Proceedings of the Sixth IJCAI,* Tokyo, pp. 300-306.

Levin, J., and Moore, J. 1077. Dialogue-games: Metacommunication structures for natural language interaction. *Cognitive Science,* vol. 1, pp. 305-420.

Levy, D. 1070. Communicative goab and strategies: Between discourse and syntax. In Givon, T., ed., *Discourse and Syntax (Syntax and Semantics,* vol. 12), Academic Press, New York, pp. 183-210.

McGuire, R., Birnbaum, L., and Flowers, M. 1081. Opportunistic processing in arguments. *Proceedings of the Seventh IJCAI,* Vancouver, B.C., pp. 58-60.

Meehan, J. 1070. *The Metanovel: Writing Stories by Computer.* Garland Publishing, New York.

Reicman, R. 1081. Modeling informal debates. *Proceedings of the Seventh IJCAI,* Vancouver, B.C., pp. 19-24.

Sacerdoti, £. 1077. *A Structure for Plans' and Behavior.* Elsevier North-Holland, New York.

Schank, R. 1077. Rules and topics in conversation. *Cognitive Science,* vol. 1, pp. 421-441.

Wilensky, R. 1083. *Planning and Understanding.* Addison-Wesley, Reading, MA.