Deciding to Remind During Collaborative Problem Solving: Empirical Evidence for Agent Strategies

Pamela W. Jordan
Intelligent Systems Program
University of Pittsburgh
Pittsburgh PA 15260
jordan@isp.pitt.edu

Marilyn A. Walker
ATT Research Laboratories*
600 Mountain Ave
Murray Hill NJ 07974
walker@research.att.com

Abstract

Previous work suggests that reminding a conversational partner of mutually known information depends on the conversants' attentional state, their resource limits and the resource demands of the task. In this paper, we propose and evaluate several models of how an agent decides whether or not to communicate a reminder. We elaborate on previous findings by exploring how attentional state and resource bounds are incorporated into the decision making process so that reminders aid the performance of agents during collaborative problem solving. We test two main hypotheses using a multi-agent problem solving simulation testbed: (1) an agent decides to present salient knowledge only when it reduces overall problem solving effort (2) an agent can use its own attentional state as a model of the attentional state of its partner when assessing the effort trade-offs of communicating a reminder. Our results support both hypotheses, suggesting that the models we propose should be further tested for multi-agent communication in problem solving situations.

Introduction

Recent work in multi-agent communication has begun to address the problem of information overload and selective attention: the problem of how an agent decides which subset of a potentially large set of facts should be attended to (Joshi 1978; Walker & Rambow 1994; Giunchiglia *et al.* 1993) *inter alia*. In human-human communication, both agents recognize that agents have selective attention and will sometimes REMIND the other agent of mutually known facts that should be selectively attended to. We posit that models of reminding for human-human conversation may be applicable to systems for both human-agent and agent-agent communication. Thus, given a situation of two communicating agents, agent A and agent B, we draw from analyses of human-human communication in order to form hypotheses of how agent A **decides** to remind agent B of a particular relevant fact. We then test our hypotheses using a dialogue simulation testbed.

In previous work on reminding, Walker and Rambow (henceforth W&R) discuss the following excerpt of a natural dialogue about which route to take (Walker & Rambow 1994):

(1)A-1: Let's walk down Walnut St.

A-2: It's shorter.

In the context of this dialogue, agent A said (A-2) despite the fact that it was already known to the hearer, agent B. Contrary to what we see in (1), a common assumption of theories of communicative behavior has been that A should not say information that B already knows (Allen 1983; Grice 1967). W&R argue that A says (A-2) in order to motivate the hearer, B, to accept the proposal in (A-1). Furthermore, they claim that A says (A-2) because she believed that although B **knew** the fact realized by (A-2), B was not **attending** to that fact.

We adopt their terminology and call (A-2) a WARRANT for the proposal. WARRANTS are beliefs used in deliberation when deciding whether to accept or reject a proposal. We also use their term SALIENT to describe the subset of an agent's beliefs and intentions that are currently attended to.[1]

W&R argue that the problem with previous models is that they do not take into account: (1) the processing involved in making relevant beliefs salient; and (2) the processing involved in making inferences, which depends on salient beliefs. They define two communicative strategies: one called Explicit-Warrant in which an agent **always** includes the warrant with a proposal and another called All-Implicit in which an agent **never** includes the warrant for a proposal. They empirically evaluate the tradeoffs between using these two strategies by simulating dialogues and then evaluating their effectiveness. They report a number of results, to be discussed below, using a performance measure that we also use below.

However, W&R simply parametrized agents for particular communicative strategies and examined when these strategies improved performance. They did not define or test a decision algorithm by which two communicating agents could decide on-line whether or not to include the warrant while they are engaged in collaborative problem solving. The goal of this paper is to define and test such a decision algorithm.

*This work was partially funded by NSF grant IRI-9314961 at the University of Pittsburgh, Mitsubishi Electric Research Laboratories, ARO Grant DAAL03-89-C0031PRI and DARPA grant N00014-90-J-1863 at the University of Pennsylvania, and Hewlett Packard Laboratories, U.K.

[1]Salient is a cognitive term, but facts in a cache can be viewed as salient facts. The critical assumption is that the salient set is a subset of what is known that is being selectively attended to.

The plan for the paper is as follows. First we present three hypotheses about the basis of an agent's algorithm for deciding whether to remind. Next we describe the Design-World dialogue simulation testbed that we use for testing our hypotheses. Finally, we present our results, which support our two main hypotheses.

Deciding to Remind in Collaborative Problem Solving

In this work, we adopt W&R's definition of performance evaluation in collaborative problem solving. This performance measure assumes that the agents are working together as a **team** (Levesque, Cohen, & Nunes 1990; Grosz & Sidner 1990), and as a team are attempting to maximize performance. PERFORMANCE is the difference between an objective measure of the utility of a completed task and a cost measure called COLLABORATIVE EFFORT (Clark & Schaefer 1989; Brennan 1990; Zukerman & McConachy 1993).

$$\text{PERFORMANCE} = \text{Task Defined RAW SCORE} - \text{COLLABORATIVE EFFORT}.$$

COLLABORATIVE EFFORT consists of all the processing required for both agents to carry out the problem solving task. In our processing architecture, this is composed of COMMUNICATIVE EFFORT, RETRIEVAL EFFORT and INFERENCE EFFORT. To ensure that our calculations of collaborative effort are independent of the implementation, the calculation of COMMUNICATIVE EFFORT is parameterized by COMMCOST, which specifies the cost of producing and understanding one message, the calculation of RETRIEVAL EFFORT is parameterized by RETCOST, which specifies the cost of one retrieval from memory, and the calculation of INFERENCE EFFORT is parameterized by INFCOST, which specifies the cost of making one inference, as defined below:

$$\begin{aligned}&\text{COLLABORATIVE EFFORT} =\\&(\text{COMMCOST} \times \text{total messages for both agents})\\&+ (\text{RETCOST} \times \text{total retrievals for both agents})\\&+ (\text{INFCOST} \times \text{total inferences for both agents})\end{aligned}$$

Our experiments focus on decisions that have to do with **reminding** another agent of mutually known information. Since at least one agent is already attending to this information, reminding cannot increase RAW SCORE; it can only reduce COLLABORATIVE EFFORT. Thus, our first hypothesis is what we call the PERFORMANCE hypothesis:

> PERFORMANCE hypothesis: Agent A decides to present salient knowledge to agent B only when it improves performance by reducing collaborative effort.

In particular, agent A's decision model must determine whether or not performance will be improved if B is reminded of the warrant for a proposal. When we attempt to define a decision algorithm to test the PERFORMANCE hypothesis, the first question that arises is how agent A goes about calculating collaborative effort. Agent A must have access to two types of information.

First, the agent must be able to access an estimate of the costs of the various processes that contribute to collaborative effort. To provide an agent with cost estimates for various processes, we define COMMCOST and RETCOST as parameters of the environment that an agent has access to. Since our decision algorithms for reminding do not involve inference effort, it is ignored by the agents, and INFCOST is set to 0.

Second, agent A must have a model of agent B's attentional state to determine whether a warrant is retrievable, and to estimate total RETRIEVAL EFFORT. Where does A get such a model? One possibility proposed by W&R (Walker & Rambow 1994) is that agent A maintains a detailed model of agent B's attentional state. This approach is consistent with work on modeling the other agent's cognitive or attentional state (Zukerman & McConachy 1993; Grosz & Sidner 1986). This possibility is the ESP hypothesis:

> ESP hypothesis: Agent A maintains a detailed model of agent B's attentional state.

We propose that another possible source of an attentional model is for agent A to use **its own** attentional state to approximate which of B's beliefs are salient, and to estimate total RETRIEVAL EFFORT. We call this possibility the SOLIPSISTIC hypothesis:

> SOLIPSISTIC hypothesis: Agent A uses its own attentional state as a model of B's attentional state.

The SOLIPSISTIC hypothesis is plausible because agent A is not always in a position to evaluate what agent B is currently attending to, and because agents in conversation appear to expend a great deal of effort to stay coordinated (Brennan 1990; Thomason 1990). In addition, in our analyses of human-human dialogues we found evidence that humans use an approximate model of one another since they sometimes make the wrong decision.[2] In the problem solving dialogue in (2), in which two human agents must negotiate the floor plan for a two room house, speaker J chooses not to fully motivate the proposal in (J-2).

(2)D-1: The green rug looks good, let's go ahead and get it. I only have $50 left – how about two yellow chairs for $25 each?

J-1: Are the yellow chairs for the living room?

D-2: Yes.

J-2: Okay, that's fine then. How about a floor lamp for 50 dollars for the living room?

D-3: Ok, we could use some light – what color is it?

J-3: Yellow, of course! I wouldn't dream of putting a red or green lamp with yellow chairs.

In dialogue (2), J and D mutually know they have a goal to match the colors of items in a room, and J assumes this goal is salient. However the color matching goal must not be salient for D since D asks about the color of the proposed item in (D-3). A cognitively plausible cause of this error is that J based her decision not to communicate the warrant on her own attentional state.

[2]Contact Rich Thomason thomason@isp.pitt.edu for information about the corpus.

The ESP and SOLIPSISTIC hypotheses give rise to two different bases for adaptive communicative strategies, depending on where agent A gets its model of agent B's attentional state. Given these adaptive communicative strategies, we can then test the PERFORMANCE hypothesis by determining whether either of these strategies improves performance when compared to the static All-Implicit and Explicit-Warrant strategies. In addition, we can test whether it is worthwhile maintaining a detailed model by comparing the performance of the two adaptive strategies.

The Design-World Testbed

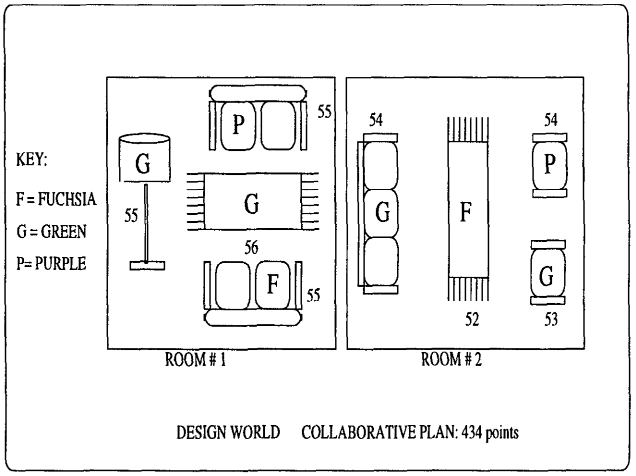

Figure 1: One Final State for Design-World Standard Task: Represents the Collaborative Plan Achieved by the Dialogue, 434 points

We test the hypotheses described above using the Design-World dialogue simulation testbed. Design-World is a testbed for theories about collaborative problem solving dialogues, that supports experiments on the interaction of agents' resource limits, their communicative strategies and the complexity of the problem solving task.[3] This section is drawn directly from the description of the Design-World testbed in (Walker & Rambow 1994).

The Design-World task requires two agents to carry out a dialogue in order to negotiate an agreement on the design of the floor plan of a two room house (Whittaker, Geelhoed, & Robinson 1993). Figure 1 shows a potential final design plan negotiated via a (simulated) dialogue, such as that in (3). Both the artificial language that the agents communicate with and an *italicized* gloss generated from that language are shown:[4]

(3)A-1: *Let's put the green rug in the study.*
(propose agent-A agent-B option-10: put-act (agent-A green rug room-1))

B-1: *Then, let's put the green lamp there.*
(propose agent-B agent-A option-33: put-act (agent-B green lamp room-1))

A-2: *Next, let's put the green couch in the study.*
(propose agent-A agent-B option-45: put-act (agent-A green couch room-1))

B-2: *No, instead let's put the purple couch in the study.*
(reject agent-B agent-A option-56: put-act (agent-B purple couch room-1))

The remainder of this section describes the initial state of the simulation, the model of dialogue, and the processing involved in generating the dialogue simulation in (3).

In each dialogue simulation, the two agents are homogeneous except for differing initial beliefs. The agent architecture is based on the IRMA architecture for resource-bounded agents (Bratman, Israel, & Pollack 1988; Pollack & Ringuette 1990), which models the planning and deliberation aspects of agent problem-solving. At the beginning of the simulation, both agents know the structure of the DESIGN-HOUSE plan. Each agent is informed of 12 items of furniture that it can use in the plan, for a total of 24 items that can be considered during planning. The 12 items assigned to an agent are stored as the agent's beliefs and neither agent is aware of what items were assigned to its partner.

Each agent is also informed of the utility of each of the 24 items that can be used in the plan and this information is stored as the agent's beliefs as well.[5] The propositions about the utility of furniture items are WARRANTS for supporting deliberation similar to the naturally occurring (A-2) in dialogue (1). The items used in the plan are assigned a utility ranging from 10 to 56. Assigning utility serves two functions: (1) it is the basis for the agent's deliberation about which plan options are better; and (2) it provides the basis for an objective performance evaluation metric of the agents' communicative behaviors.

Because attentional state is an important factor in communicative effectiveness, the Design-World version of IRMA includes a model of limited Attention/Working Memory (AWM) that is similar to a cache and that is based on human working memory (Landauer 1975; Walker 1996). AWM defines those salient beliefs that an agent selectively attends to. Due to the way AWM is modelled, the salience of all the information used in reasoning and planning varies according to the recency and frequency with which it is accessed. Furthermore, the size of AWM is parameterized so that we can test the relationship between communicative strategies such as reminding and the degree to which the size of the subset of beliefs that an agent can selectively attend to is limited. We test 4 different levels of attention limitation in the experiments below.

[3]We will not attempt to convince the reader of the model of collaborative planning implemented in Design-World since that is done elsewhere (Walker & Rambow 1994). We merely summarize Design-World for the convenience of the reader and assume that it is a useful testbed for validating our hypotheses. Contact Marilyn Walker for the simulation software.

[4]The generation of the gloss was not a focus of Design-World and is done automatically by adhoc methods.

[5]The agents don't know that all items have been assigned and don't engage in existential reasoning, i.e. they do not make inferences about what better choices other agents might propose on the basis of these utilities.

Modelling Collaborative Interactions

In each dialogue simulation, the agents' goal is to **agree on** 4 pieces of furniture for each room so that the maximum utility is achieved. Negotiating an agreement between two agents consists of a cycle of four steps. First, individual agents perform means-end reasoning about OPTIONS to use various furniture pieces that they have. Second, individual agents deliberate about which options are preferable based on the utility of using the item in the plan. Since the AWM parameter limits what an agent is attending to, an agent may not be able to recall and identify its highest utility pieces at any given time.

Once a preferred option is identified, the agent attempts to get the other agent to agree to a proposal involving that option. In general, agents' communicative intentions are realized through DISCOURSE ACTS such as PROPOSALS, ACCEPTANCES, REJECTIONS and CLARIFICATIONS. In the third step, agents use the preferred options to make PROPOSALS to other agents to PUT a piece of furniture into a room. Then in the fourth step, the other agents ACCEPT, REJECT, or request CLARIFICATION of these proposals.

This is illustrated by the dialogue excerpt in (3). After receiving utterance (A-1) from agent A, agent B conducts means-end reasoning about the plan-step that A has made a proposal about. It then evaluates A's proposal by comparing it with the options it has generated by reasoning, and on the basis of this comparison (deliberation), it decides whether to accept or reject it. During this evaluation, it will attempt to retrieve the warrant propositions stored earlier in memory which are the beliefs that allow it to evaluate each proposal and to compare another agent's proposal with the options that it has generated by its own means-end reasoning. Remember that B knows the utility information for all the items that he and A could propose, but that information may not be salient.

Proposals, such as (A-1) and (B-1) in (3), are inferred to be implicitly ACCEPTED because they are not rejected (Carberry 1989). If a proposal is ACCEPTED, either implicitly or explicitly, then the option contained in the proposal becomes a mutual intention that contributes to the final design plan (Power 1984; Sidner 1994). Agents REJECT a proposal if deliberation produces an option with a higher utility. For example, in (B-2) B rejects the proposal of option-45 in (A-2), proposing instead option-56. Either B could not recall the utility of option-45, or the utility of option-56 is higher.

Discourse Strategies for Hypothesis Testing

A discourse strategy is a particular way of achieving a DISCOURSE ACT such as a PROPOSAL. Agents are parametrized for different discourse strategies by placing different expansions of discourse acts in their plan libraries. To test the hypotheses discussed earlier, we examine four strategies: (1) All-Implicit (2) Explicit-Warrant (3) Solipsistic (4) ESP. Each strategy varies the decision algorithm that is used when an agent is reasoning about whether to expand a proposal to include a reminder of the warrant for the proposal.

The All-Implicit strategy **never** includes warrants in proposals, leaving it up to the other agent to retrieve them from memory. An agent utilizing this strategy acts as though it assumes that everything the other agent knows is salient. In dialogue (3), both agents communicate using the All-Implicit strategy, and PROPOSALS expand to the PROPOSE communicative acts shown in (A-1), (B-1), and (A-2).

The Explicit-Warrant strategy **always** expands the PROPOSAL discourse act to be a WARRANT followed by a PROPOSE utterance (Suthers 1993). An agent utilizing this strategy acts as though the other agent never retains anything in working memory, as though nothing is salient for the other agent. For example in (4) (A-1) is agent A's WARRANT for his proposal in (A-2):

(4)A-1: **Putting in the green rug is worth 56.**

A-2: *So, let's put the green rug in the study.*

B-1: **Putting in the green lamp is worth 55.**

B-2: *So, let's put the green lamp in the study.*

The final two strategies **sometimes** include a warrant depending on the source of information that agent A uses about B's attentional state in its decision algorithm. These are the strategies based on the ESP and SOLIPSISTIC hypotheses discussed earlier; we call these the Solipsistic strategy and the ESP strategy.

In the ESP strategy, agent A maintains a detailed model of agent B's attentional state. To implement the ESP strategy, we actually give agent A access to agent B's mind in the simulation. ESP provides data on how a perfect attentional model of the other agent affects performance.

An agent using the Solipsistic strategy uses its own attentional state as a model for the other agent. To implement the Solipsistic strategy, we endow agent A with the capability of keeping track of its own cognitive effort for all of the processing that A does. Agent A then uses its own retrieval effort to determine whether it improves performance to remind B of a warrant for a proposal rather than letting B retrieve the warrant from B's own memory.

Evaluating Performance

Earlier, we discussed the model of performance that we assume, in which agents work as a **team**. Performance is the difference between the RAW SCORE for the task and COLLABORATIVE EFFORT. RAW SCORE is task specific: we simply sum the utility of the furniture pieces in each PUT action in the final plan. For example, the raw score for the design plan in Figure 1 is 434 points.

As we discussed above, COLLABORATIVE EFFORT is composed of COMMUNICATIVE EFFORT and RETRIEVAL EFFORT. Our PERFORMANCE hypothesis is that the adaptive reminding strategies can improve performance by reducing collaborative effort. Thus, the decision algorithm for both the ESP and Solipsistic strategies is to say the warrant whenever the model of agent B's AWM predicts that (1) the warrant is not salient or (2) the warrant will be more costly for agent B to retrieve than for agent A to communicate.

Note however, that any savings in retrieval effort is always at the cost of an additional utterance, the reminder, which increases communicative effort. To provide a fair test of our adaptive strategies, we construct two adversarial dialogue situations, one which favors reminding by making each retrieval relatively costly, and one which favors not reminding by making each communicated message relatively costly. We do this by using two different parameter settings for COMMCOST and RETCOST in our experiments when calculating collaborative effort. Retrieval effort dominates with unit cost settings of COMMCOST = .1 and RETCOST = .001. Communicative effort dominates with unit cost settings of COMMCOST = .001 and RETCOST = 1×10^{-6}.

We experiment with four AWM ranges for the resource limits; LOW for very resource limited agents, MID for agents hypothesized to be similar to human agents, MIDH for agents less constrained than humans, and HIGH for resource unlimited agents. An AWM range indicates the maximum number of beliefs and intentions that are potentially salient for an agent at any processing step.[6]

We determined that a sample size of 200 dialogues per experimental condition is adequate for determining whether a strategy affects PERFORMANCE. To collect these samples we simulate 200 dialogues with the appropriate parameter settings, yielding a performance distribution for each strategy and set of assumptions tested.

We present our comparisons of dialogue strategies in plots of the **differences** in the mean PERFORMANCE of agent pairs using different strategies, such as the graph in Figure 2. In Figure 2, agent pairs using the ESP strategy are compared against agent pairs using the All-Implicit strategy. **Differences** in the mean PERFORMANCE between the two strategies are plotted on the Y-axis against the four AWM parameter ranges on the X-axis. Each point in the graph at a particular AWM range is the difference in the mean performances of 200 samples where the agent pairs both use the ESP strategy and 200 samples where the agents use the All-Implicit strategy. This graph summarizes the information from 8 total performance distributions (1600 simulated dialogues).

To see which of the performance differences are significant we run planned comparisons using one-way analysis of variance (anova) where statistical significance is determined by the modified Bonferroni test (MB). When comparing two strategies for the same AWM range and unit cost settings, if the mean performance of strategy 1 is significantly less than the mean of strategy 2, according to the MB test, then strategy 1 is DETRIMENTAL and strategy 2 is BENEFICIAL for agents in that AWM range.

Experimental Results

In initial experiments, we duplicated W&R's results. They found that the Explicit-Warrant strategy (always say the warrant) compared to the All-Implicit strategy (never say the warrant) is beneficial at the two highest AWM ranges, when retrieval effort dominates other processing effort, but that it is detrimental at the MID AWM range when communication effort dominates, except when the task requires agents to be coordinated on warrants underlying intended actions (Walker & Rambow 1994).

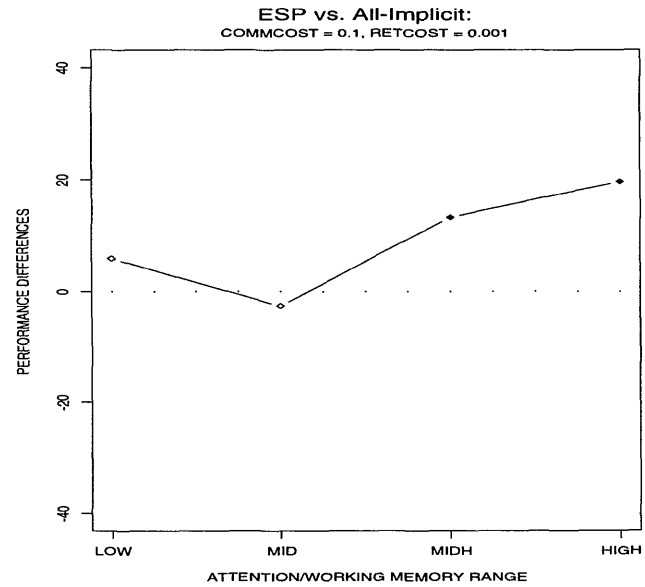

Figure 2: When retrieval effort dominates, ESP is beneficial compared to All-implicit at MIDH and HIGH AWM ranges. The other points are not statistically significant.

Now, our main goal is to test our adaptive strategies. First, we compare the performance of the two adaptive discourse strategies to the static Explicit-Warrant and All-Implicit strategies, in both adversarial dialogue situations. Second, we compare the two adaptive strategies in both situations to see whether any significant performance differences arise.

In order for an adaptive strategy to be worthwhile, it should perform at least as well as the static strategy that is best in a particular dialogue situation. Furthermore, we expect the adaptive strategies to be advantageous by improving performance in **all** dialogue situations, while static strategies only perform well in **some** dialogue situations.

Thus, given W&R's results, we expect the adaptive strategies to perform as well as All-Implicit when communication effort dominates. This means that (1) when an adaptive strategy is compared to the All-Implicit strategy, it should be neither detrimental nor beneficial, and (2) when an adaptive strategy is compared to the Explicit-Warrant strategy, it should be beneficial at MID AWM.

Similarly we expect the adaptive strategies to perform as well as Explicit-Warrant when retrieval effort dominates. This means that (1) when an adaptive strategy is compared to Explicit-Warrant, it should be neither detrimental nor beneficial; and (2) when an adaptive strategy is compared to All-Implicit, it should be beneficial at MIDH and HIGH AWM ranges.

[6]Because the AWM model is probabilistic, each dialogue simulation for an AWM range has a different result. Also as the maximum number of salient beliefs and intentions increases (as we increase the AWM range parameter), the number of retrievals during a dialogue also increases since more of memory is available to search. In addition, when more memory is available to search, more messages are exchanged since more options are available to discuss.

	Exp-Warr vs. All-Imp			Adapt vs. All-Imp			Adapt vs. Exp-Warr		
	MID	MIDH	HIGH	MID	MIDH	HIGH	MID	MIDH	HIGH
Retrieval Dominates	-	Exp-Warr	Exp-Warr	-	Adapt	Adapt	-	-	-
Communication. Dominates	All-Imp	-	-	-	-	-	Adapt	-	-

Table 1: Expected beneficial AWM ranges for the Adaptive Strategies ("-" indicates no performance difference)

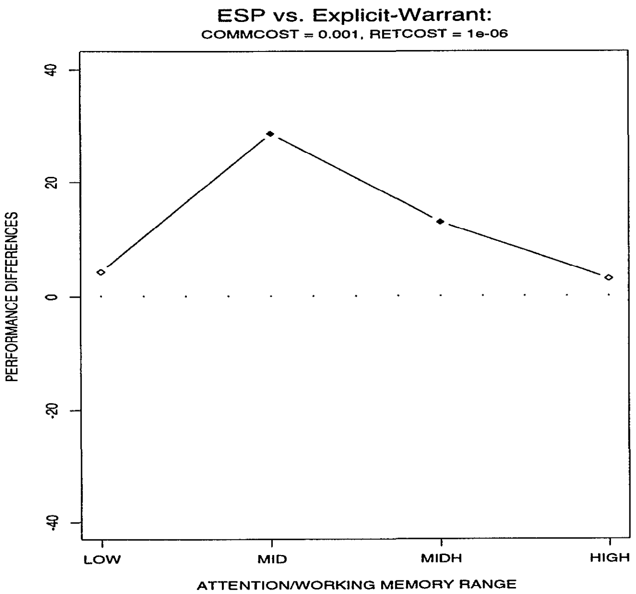

Figure 3: When communication effort dominates, ESP is beneficial compared to Explicit-Warrant at MID, and MIDH AWM ranges.

Table 1 summarizes the behavior of Explicit-Warrant vs. All-Implicit and the expected behaviors for the adaptive strategies. As Table 1 shows, we predict that the adaptive strategies will perform better under more conditions than either of the two static strategies. If our predictions in Table 1 are correct, we should see significant performance differences for only Adaptive vs. Explicit-Warrant when communication effort dominates and for only Adaptive vs. All-Implicit when retrieval effort dominates. In the discussion of the results, we include difference plots only for results with significant performance differences.

ESP vs. All-Implicit and Explicit-Warrant

To test the PERFORMANCE hypothesis we factor out the source of the attentional state by using the ESP model. Using this perfect model enables us to determine what happens when the calculation of collaborative effort is completely accurate.

When retrieval effort dominates the calculation of collaborative effort, the ESP strategy is beneficial compared to All-Implicit at the MIDH and HIGH AWM ranges as shown in Figure 2 (MB(MIDH) = 8.14, $p < .01$ and MB(HIGH) = 17.77, $p < .002$). As Table 1 shows, this is what we predicted.

When communication effort dominates the calculation of collaborative effort, the ESP strategy is beneficial compared to Explicit-Warrant at MID and MIDH as shown in Figure 3 (MB(MID) = 14.52, $p < .002$ and MB(MIDH) = 8.04, $p < .01$). However our expectation was that it would only be beneficial at the MID range. The unexpected benefits of the ESP strategy at the MIDH range of AWM are perhaps due to the combined savings in retrievals and communications being significant.

The results of the experiments with the ESP strategy, exceed the expectations in Table 1. But since our minimum expectations were met for what would happen if the PERFORMANCE hypothesis is true, there is positive evidence for this hypothesis. The decision to present a salient warrant when it reduces collaborative effort does increase overall performance.

Solipsistic vs. All-Implicit and Explicit-Warrant

To test the SOLIPSISTIC hypothesis, we look at the results of the Solipsistic strategy where agent A uses its own processing to estimate B's effort and uses this to calculate collaborative effort. The precision of the model is what we wish to test. There are two ways in which the Solipsistic model can fail to be a good approximation in the collaborative effort assessment: (1) the warrant may not be salient for agent B (2) the ease of accessing a salient warrant may be different for agent A and agent B.

Despite these potential failings, the Solipsistic strategy met the predictions shown in Table 1 exactly. When retrieval effort dominates the calculation of collaborative effort, Solipsistic is beneficial compared to the All-Implicit strategy at the MIDH and HIGH AWM ranges as shown in Figure 4 (MB(MIDH) = 4.51, $p < .05$ and MB(HIGH) = 36.14, $p < .002$). When communication effort dominates, Solipsistic is beneficial compared to Explicit-Warrant at the MID AWM range as shown in Figure 5 (MB(MID) = 4.34, $p < .05$).

Since our expectations were met for what would happen if the SOLIPSISTIC hypothesis were true, there is positive evidence for this hypothesis. Using agent A's own attentional state to model agent B is accurate enough to improve the performance of the adaptive strategy, despite the fact that the initial state of the agents is different and despite the fact that our AWM model is a probabilistic one.

ESP vs. Solipsistic

Finally, since the ESP and SOLIPSISTIC hypotheses are mutually exclusive, we compare the two adaptive strategies for performance differences. These experiments show that ESP compared to Solipsistic is beneficial at only the HIGH AWM range when communication effort dominates (MB(HIGH) = 8.31, $p < .01$) (no figure due to lack of space).

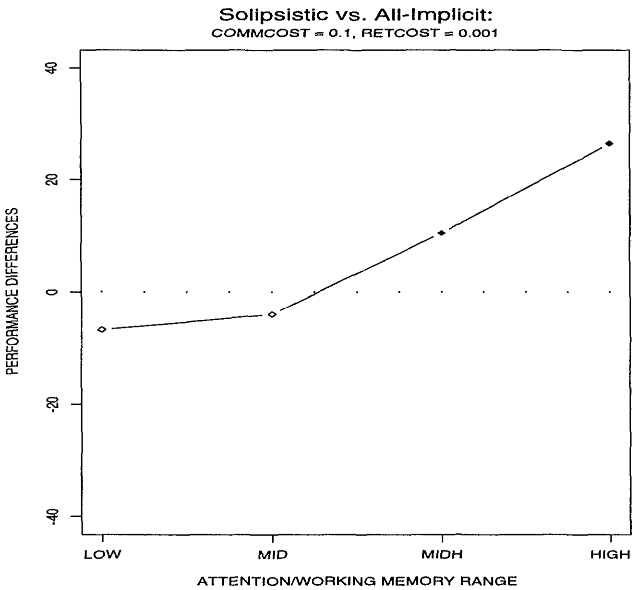

Figure 4: When retrieval effort dominates, Solipsistic is beneficial compared to All-Implicit at MIDH and HIGH AWM ranges.

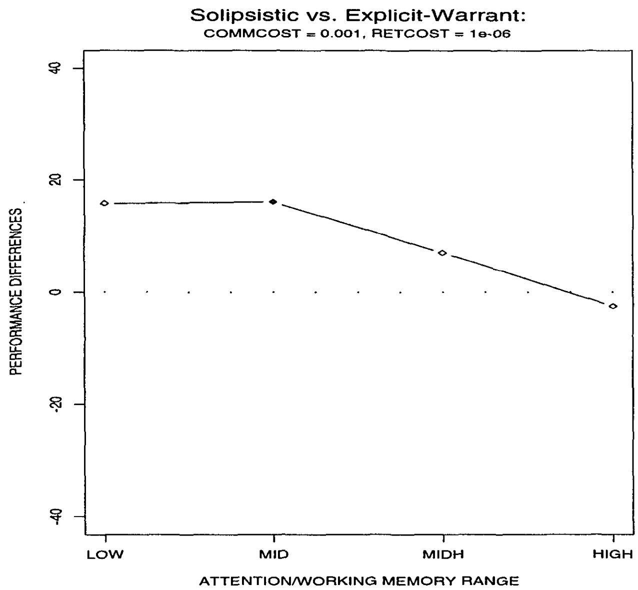

Figure 5: When communication effort dominates, Solipsistic is beneficial compared to Explicit-Warrant at the MID AWM range.

This shows that in at least some communicative situations, it may not be worth maintaining a detailed model of agent B, given the potential overhead of maintaining such a model, which we did not include in the calculation of collaborative effort.

To test this idea further, we added small percentages of the task defined RAW SCORE to collaborative effort to account for the overhead involved in maintaining the ESP model, and then compared performance again. We found that if the effort of maintaining the ESP model is greater than 4% of the RAW SCORE, that Solipsistic performs as well as ESP at all AWM settings.

Discussion

Our goal was to test three hypotheses about deciding to remind: the PERFORMANCE, SOLIPSISTIC, and ESP hypotheses. We found support for the PERFORMANCE hypothesis: the decision to present a warrant should be based on whether doing so enhances overall performance. This is not particularly surprising.

To test the PERFORMANCE hypothesis, we evaluated two adaptive strategies for reminding, Solipsistic and ESP, in comparison with similar static strategies. Furthermore, we evaluated them under two different adversarial dialogue conditions, one which favors reminding and one which favors not reminding. We found that the adaptive strategies never perform worse than the static strategies. Furthermore, both adaptive strategies are superior to the two static strategies since a single adaptive strategy performs well over all dialogue conditions.

We also found support for the SOLIPSISTIC hypothesis over the ESP hypothesis. The SOLIPSISTIC hypothesis is, that in deciding to remind, agent A can approximate B's attentional state with its own. The ESP hypothesis is that agent A maintains a detailed model of agent B's attentional state to use in making decisions about reminding. To our knowledge, no-one has ever proposed or tested the SOLIPSISTIC hypothesis before. By comparing the Solipsistic strategy to the ESP strategy, we found that it may not be worth maintaining a detailed model of agent B's attentional state, given the potential costs of doing so.

We interpret the results on the SOLIPSISTIC hypothesis to mean that the fidelity of the AWM model to the recency and frequency properties of human working memory produces an environment in which agent A's attentional state is probabilistically correct as an approximation of B's. With respect to modelling humans, this suggests that it may be a cognitively efficient strategy to use one's own attentional state to estimate that of your conversational partner. With respect to building artificial agents, the results suggest that when two agents are engaged in a synchronous collaborative problem solving dialogue, the Solipsistic model may be a useful approximation for agent A to use in deciding when to remind agent B of relevant facts.

Our methodology for achieving these results consists of a controlled experimental environment in which we can manipulate variables that are relevant to our hypotheses. The environment implements a particular formal model of collaborative problem solving. Thus, our experiments are directly affected by both the formal model and the experimental variables that we manipulate, and provide a way to evaluate how experimental variables interact with the formal model.

Hanks, Pollack and Cohen discuss at length the importance of demonstrating that results collected in simulation experiments generalize beyond the particulars of a testbed environment (Hanks, Pollack, & Cohen 1993; Cohen 1995). There are several reasons why we expect the results above to generalize. First, our simulation is based on a model of collaborative planning that is similar to other models (Grosz & Sidner 1990; Levesque, Cohen, & Nunes 1990; Guinn 1994; Cohen 1995), and based on general assumptions about the underlying agent architecture (Bratman, Israel, & Pollack 1988). Second, we tested reminding for a particular information configuration: agent A makes a proposal and reminds B of the warrant that supports it. Since this is a general information configuration that is found in any agent that deliberates, our models of reminding should extend to other problem-solving situations. Third, our AWM model has many similarities to a cache, thus Solipsistic strategies should be beneficial in any architecture with a cache type memory such a SOAR (Lehman, Lewis, & Newell 1991), as long as communication is approximately synchronous. Finally, because our performance evaluation depends on calculating performance for the team (Levesque, Cohen, & Nunes 1990; Grosz & Sidner 1990), we believe that our reminder models will generalize to any team-oriented problem-solving environment such as Phoenix (Cohen 1995).

In future work, the implicated decision model can be tested in human-computer interaction, as in intelligent tutors, to see how well the model extends to non-homogeneous agents (Moore 1994). Additional experiments could also determine whether the Solipsistic strategy breaks down under some circumstances, when the ESP hypothesis might hold, and whether adaptive strategies are beneficial compared to the static strategies under less extreme processing assumptions.

References

Allen, J. F. 1983. Recognizing intentions from natural language utterances. In Brady, M., and Berwick, R., eds., *Computational Models of Discourse*. MIT Press.

Bratman, M.; Israel, D.; and Pollack, M. 1988. Plans and resource bounded practical reasoning. *Computational Intelligence* 4:349–355.

Brennan, S. E. 1990. *Seeking and Providing Evidence for Mutual Understanding*. Ph.D. Dissertation, Stanford University Psychology Dept. Unpublished Manuscript.

Carberry, S. 1989. Plan recognition and its use in understanding dialogue. In Kobsa, A., and Wahlster, W., eds., *User Models in Dialogue Systems*. Berlin: Springer Verlag. 133–162.

Clark, H. H., and Schaefer, E. F. 1989. Contributing to discourse. *Cognitive Science* 13:259–294.

Cohen, P. R. 1995. *Empirical Methods for Artificial Intelligence*. Boston: MIT Press.

Giunchiglia, F.; Serafini, L.; Giunchiglia, E.; and Frixione, M. 1993. Non-Omniscient Belief as Context-Based Reasoning. In *IJCAI 93*, 548–554.

Grice, H. P. 1967. *William James Lectures*.

Grosz, B. J., and Sidner, C. L. 1986. Attentions, intentions and the structure of discourse. *Computational Linguistics* 12:175–204.

Grosz, B. J., and Sidner, C. L. 1990. Plans for discourse. In *Cohen, Morgan and Pollack, eds. Intentions in Communication, MIT Press*.

Guinn, C. I. 1994. *Meta-Dialogue Behaviors: Improving the efficiency of Human-Machine Dialogue*. Ph.D. Dissertation, Duke University.

Hanks, S.; Pollack, M.; and Cohen, P. 1993. Benchmarks, testbeds, controlled experimentation and the design of agent architectures. *AI Magazine*.

Joshi, A. K. 1978. Some extensions of a system for inference on partial information. In *Pattern Directed Inference Systems*. Academic Press. 241–257.

Landauer, T. K. 1975. Memory without organization: Properties of a model with random storage and undirected retrieval. *Cognitive Psychology* 495–531.

Lehman, J. F.; Lewis, R. L.; and Newell, A. 1991. Natural language comprehension in SOAR. Technical Report CMU-CS-91-117, Carnegie Mellon University.

Levesque, H. J.; Cohen, P. R.; and Nunes, J. H. T. 1990. On acting together. In *AAAI90*.

Moore, J. 1994. *Participating in Explanatory Dialogues:Interpreting and Responding to Questions in Context*. MIT Press.

Pollack, M., and Ringuette, M. 1990. Introducing the Tileworld: Experimentally Evaluating Agent Architectures. In *AAAI90*, 183–189.

Power, R. 1984. Mutual intention. *Journal for the Theory of Social Behaviour* 14.

Sidner, C. 1994. An artificial discourse language for collaborative negotiation. In *AAAI 94*, 814–820.

Suthers, D. D. 1993. Sequencing explanations to enhance communicative functionality. In *Proceedings of the 15th Annual Meeting of the Cognitive Science Society*.

Thomason, R. 1990. Propagating epistemic coordination through mutual defaults I. In Parikh, R., ed., *Proceedings of the Third Conference on Theoretical Aspects of Reasoning about Knowledge*, 29–39. Morgan Kaufmann.

Walker, M. A., and Rambow, O. 1994. The role of cognitive modeling in achieving communicative intentions. In *The 7th International Conference on Natural Language Generation*.

Walker, M. A. 1996. Limited attention and discourse structure. *Computational Linguistics* 22(1).

Whittaker, S.; Geelhoed, E.; and Robinson, E. 1993. Shared workspaces: How do they work and when are they useful? *IJMMS* 39:813–842.

Zukerman, I., and McConachy, R. 1993. Generating concise discourse that addresses a user's inferences. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, 1202–1207. Morgan Kaufmann Publishers, Inc.