

# The Institute For Research In Cognitive Science

**Discourse Strategies for Attention-  
Limited Agents**



by

**Marilyn A. Walker**

**University of Pennsylvania  
3401 Walnut Street, Suite 400C  
Philadelphia, PA 19104-6228**

**July 1993**

Site of the NSF Science and Technology Center for  
Research in Cognitive Science



# Discourse Strategies for Attention-Limited Agents

Marilyn A. Walker  
University of Pennsylvania  
Computer Science Dept.\*  
Philadelphia, PA 19104  
lyn@linc.cis.upenn.edu

## Abstract

This paper presents results on the role of limited attention and inference in dialogue and the effect of discourse strategies that interact with these limitations. These factors are explored through a dialogue simulation environment, Design-World, in which two artificial agents collaborate to perform a simple design task. The results are based on several experiments carried out in this environment. These experimental results demonstrate that discourse strategies that incorporate redundant information can be beneficial when agents have limited resources for attention and inference.

## 1 Introduction

A common rule of dialogue is not to tell people things that they already know; this is one reason to keep track of the COMMON GROUND, facts already mutually believed by conversants in dialogue. This ‘no redundancy’ rule is a basic tenet of many theories of dialogue, from linguistic[Gaz79], philosophical[Sta78] and computational perspectives[Coh78]. For example, by Stalnaker’s account, as a dialogue progresses, each assertion is ACCEPTED and then added to the common ground, unless it

---

\*Thanks to Aravind Joshi for funding me under ARO grants DAAG29-84-K-0061 and DAAL03-89-C0031PRI, DARPA grant N00014-85-K0018, and NSF grant MCS-82-19196 at the University of Pennsylvania, and to Hewlett-Packard Laboratories, U.K.

is REJECTED by one of the conversants. This addition functions to limit the set of worlds that the conversants believe possible. It is assumed that efficient agents would never produce an utterance which fails to further limit these worlds. However, consider example 1 from a radio talk show for financial advice:<sup>1</sup>

- (1) ( 3) e. .... – and I was wondering – should I continue on with the certificates or  
( 4) h. Well it’s difficult to tell because we’re so far away from any of them – but I would suggest this – if *all of these are 6 month certificates and I presume they are*  
( 5) e. *yes*  
( 6) h. *then I would like to see you start spreading some of that money around*  
( 7) e. uh hu  
( 8) h. Now in addition, how old are you?  
.  
(discussion about retirement investments)  
.  
(21) e. uh huh and  
(22) h. but as far as the certificates are concerned, I’D LIKE THEM SPREAD OUT A LITTLE BIT - THEY’RE ALL 6 MONTH CERTIFICATES

---

<sup>1</sup>I am grateful to Julia Hirschberg and Martha Pollack for originally transcribing this corpus and providing me with tapes of the original broadcast [PHW82].

- (23) e. yes
- (24) h. and I don't like putting all my eggs in one basket.....

Utterance 1-(22) consists of two clauses that both realize propositions discussed earlier in 1-(4) to (6). These propositions are reintroduced at the OPENING of the discourse segment marked by the phrase *but as far as the certificates are concerned*. These utterances are examples of INFORMATIONALLY REDUNDANT UTTERANCES, henceforth IRUs, in which an agent tells another agent something that is already mutually believed. The IRUs in 1 are a subclass of IRUs called ATTENTION IRUs, which redirect the listener's attention to the previously discussed propositions. Here and below, IRUs are shown in caps, while the ANTECEDENTS, which originally added the IRUs propositional content to the common ground, will be shown in *italics*.

Another type of IRU is shown in 2:

- (2) (1) Listen to Ramesh.
- (2) HE'S INDIAN. (DH 11/5/91)

In this case the IRU in 2-2 provides a WARRANT for doing the action given in 2-1. This content of this IRU had not been explicitly evoked in the discourse situation, but was certainly a mutual belief among the conversants.

In the discourse situation where 2 was said, the discussion involved a shared intention to instantiate *X* in the proposition *We should eat at restaurant X of type Indian*. Doing the action  $\gamma$  in  $U_1$  is meant to CONTRIBUTE to this shared intention[Pol90]. The IRU *Ramesh is Indian*, is a WARRANT for adopting the intention to do  $\gamma$ . In other words, the addressees were meant to infer that the proposition conveyed by  $U_2$  is a reason for adopting an intention to do  $\gamma$ [WJ82]. The reason for saying the IRU is to make the proposition that it conveys salient so that the audience can easily make this inference.

Both types of examples are evidence for the following DISCOURSE INFERENCE CONSTRAINT:

DISCOURSE INFERENCE CONSTRAINT : Inferences in dialogue are derived from a small set of propositions that are currently salient.

The discourse inference constraint may be a correlate of the fact that human agents have limited attentional capacity and thus only a small number of propositions can be salient[Mil56, KST82, Sol92]. The focus of the constraint is the relation of limited attention to inference. Other accounts of limited reasoning have also proposed that agents only focus on a subset of their knowledge at any one time[Jos78, JW81, JWW84, Lan87], and the role of this subset of beliefs is substantially the same in these lines of research: a reduction in computation that results from reasoning over a smaller set of facts. This paper argues that IRUs, as in 1 and 2 above, manipulate agents' attention and interact with their inferential capacities; thus they are a reflection of agent's resource bounds.

While the constraint as stated is still underspecified, a model of limited attention will be proposed in section 3 and section 5 will report experimental results of agents in task-oriented dialogues who can only make inferences on those facts that are currently attended to. The Design-World simulation environment is parameterizable and thus allows us to explore the interaction of attention and inference in a systematic way. The experiments reported vary agents' discourse strategies and attentional capacity and measure the effect that this has on agents' scores for the task, as well as the number of inferences and the number of memory retrieval operations required to achieve these scores.

The discourse strategies implemented in Design-World are based on the distributional analysis of IRUs in the financial advice corpus, which also provides support for the discourse inference constraint[Wal92]. Section 2 describes the Design-World environment.

The model of limited attention in Design-World will be discussed in section 3. The discourse strategies used to test the interaction of attention and inference will be discussed in section 4. Section 5 present the experimental results that support the relationship between IRUs, limits on agents' attentional capacity, and limits on inference.

## 2 Design-World

The Design-World simulation environment consists of the floor plan for a house[WGR93]. The Design-World task consists of two cooperating agents carrying out a mixed-initiative dialogue in order to reach an agreement on a design for a floor plan with two rooms. A potential Design-World initial state is shown in figure 1.




Figure 1: Initial State for the Design-World Task

A design for a room consists of any four pieces of furniture. Furniture items are of 5 types: couch, table, chair, lamp and rug. Each furniture item has a color and point value. Agents share beliefs about which pieces of furniture exist and what their colors and point values are. The only domain actions are:

- (Put ?Agent ?Furniture ?Room ?Time)
- (Remove ?Agent ?Furniture ?Room ?Time)

The goal is to design a house by designing the two rooms that make up the house. Dialogues are always between two agents A and B, and actions can be performed by either agent or by the agents as a pair, A&B. The ?Time variable for each action and state is either NOW or LATER.

The Design-World architecture is based on the IRMA architecture for resource-bounded agents[BIP88, PR90]. Design-World adds a model of limited attention to IRMA. The Design-World version of the IRMA architecture is shown in figure 2. As shown, the model of limited attention constrains the belief and means-end reasoner. This will be discussed more fully in section 3.

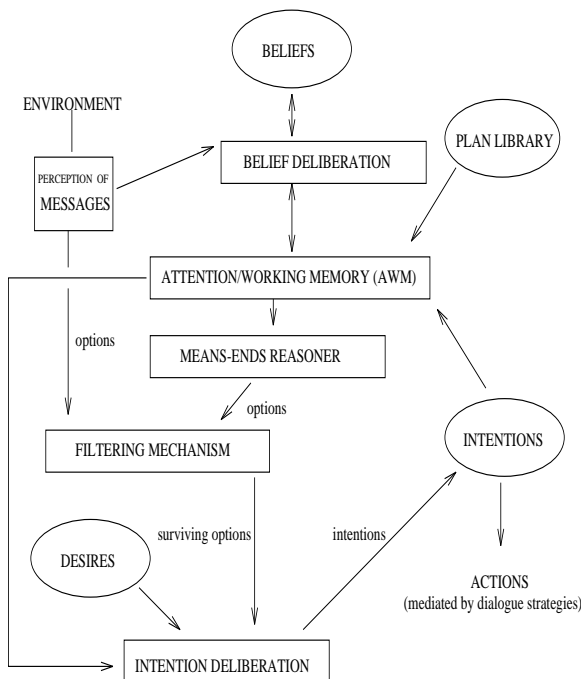


Figure 2: Design-World IRMA Agent Architecture for Resource-Bounded Agents

The agents' goal in carrying out the dialogue is to achieve a COLLABORATIVE PLAN, a collection of mutual beliefs and intentions[Pol90]. A collaborative plan is achieved by a cycle in which: (1) individual agents perform means-end reasoning about options in the domain; (2) individual agents deliberate about which options are preferable;

(3) then agents communicate **PROPOSALS** to other agents, based on the options identified in a reasoning cycle, about actions that contribute to the satisfaction of their goals; (4) then these proposals are **ACCEPTED** or **REJECTED** by the other agent.

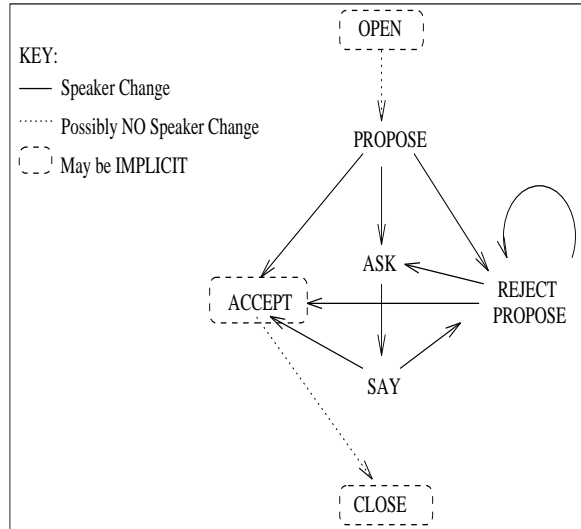


Figure 3: Dialogue Actions for the Design-World Task

A schemata for agents’ messages is given in figure 3.<sup>2</sup> Every dialogue segment about a particular subgoal of the task consists of the actions shown in figure 3. The dialogues are mixed-initiative since both agents reason about the domain and both agents have goals and either agent can make a **PROPOSAL**. As messages are sent and received during a dialogue, agents construct a corresponding structure for mutual beliefs and intentions. However, not every dialogue action shown in figure need be explicitly realized. In particular **OPEN SEGMENT**, **CLOSE SEGMENT** and **ACCEPT PROPOSAL** messages may be implicit. When they are left implicit, they must be inferred to have occurred by the other agent.

<sup>2</sup>These actions do not represent a claim that this is how humans conceptualize dialogue or that problem solving dialogues can be analyzed into structures exactly like these. The effects of resource-bounds will operate independently of the range of utterance level intentions and other strategies available to an agent. See [Car92], [CGRJ92], and [Sid92] for other sets of utterance intentions for dialogues.

These variations, in what is left implicit and what is made explicit, as well as the form of a proposal, acceptance etc., are determined by agents’ discourse strategies, which will be discussed in section 4. Also, as shown, rejections always include a counter-proposal because the decision to reject a proposal is based on comparing the proposal with other options an agent knows about, and another option must be of greater utility to support a rejection. Acceptance/rejection can also be postponed by **ASKING** for more information, and as a simplifying assumption, interaction is limited so that B will always ask for more information if B doesn’t have enough information to make a comparison between the options that B knows about and the one that A has proposed. However, if B has no options available to satisfy the same goal, B can accept A’s proposal without knowing its exact value.

Every segment results in an acceptance because, since agents share a global utility function based on the points accumulated for a particular design, all conflicts are eventually resolved. The acceptance of a proposal generates the mutual beliefs necessary for the agents to have a **COLLABORATIVE PLAN**. This is reflected in the definition given below (See also [WS88, WW90, CL91, GS90, Sid92]):

(**COLLABORATIVE-PLAN** A&B  
(Achieve A&B Goal))  $\iff$

1. (MB A&B (Intend A  $\vee$  B,  $\alpha_1 \wedge \dots \alpha_n$ ))
2.  $\forall i$  (MB A&B (Contribute  $\alpha_i$  (Achieve A&B Goal)))
3. (MB A&B (Max-Utility  $\alpha_1 \wedge \dots \alpha_n$ ) (Achieve A&B Goal))

Figure 4 shows a potential final collaborative plan for the Design-World task as a result of the dialogue. The points associated with a furniture item supports the calculation of utility of including that item in the design plan. The agents’ decisions of whether

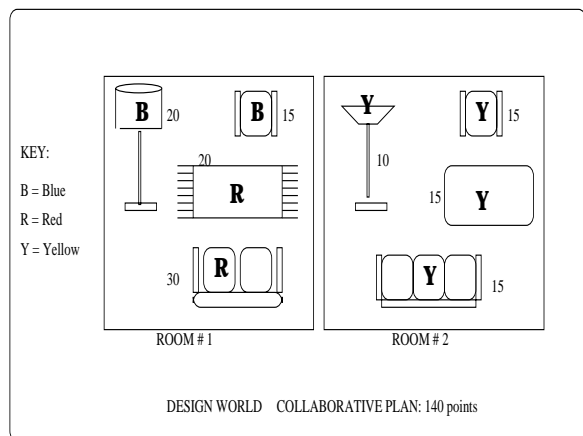


Figure 4: Potential Final State for Design-World Task: Represents the Collaborative Plan Achieved by the Dialogue

to make, accept or reject proposals are based on deliberating about which future actions will maximize the points achieved by their design. Associating utility with the final plan means that the effectiveness of different dialogue strategies can be measured. Overall performance is measured by the points associated with the final design as compared with the costs to achieve these points. Costs are based on: (1) number of utterances in the dialogue; (2) the amount of means-end reasoning required; and (3) the number of steps involved in searching the representation of current beliefs and intentions. Thus it is possible to explore in a principled manner the trade-offs associated with different dialogue strategies.

### 3 Attention/Working Memory Model

The attention/working memory model, AWM, used in Design-World is a very simple model that incorporates both recency and frequency effects, and which also accounts for a broad range of empirical results on human memory and learning[Lan75]. This model consists of a three dimensional space in which items are stored in chronological sequence according to the current location of

a moving memory pointer. The next memory pointer location is within two steps of the current location, however, of the distance 2 locations, the location to be used next is randomly selected. This is depicted in figure 5. The current memory pointer is the memory location at the end of the path shown.

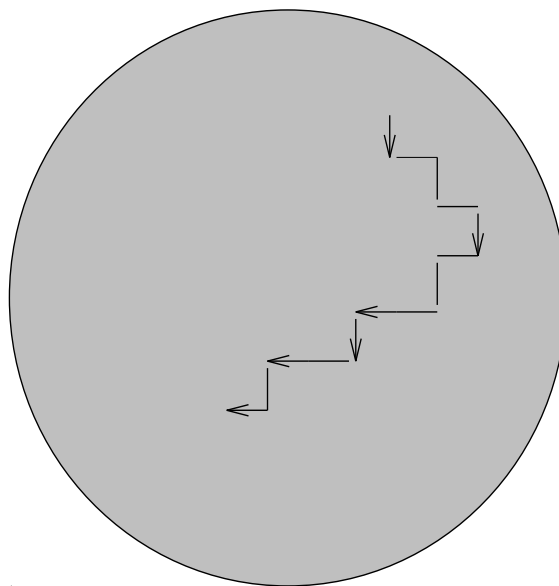


Figure 5: Memory Structure: Three Dimensional Store, Random Storage

Thus memory items based on experiences that occurred near one another in time will be closer together, but there is a random choice associated with the exact memory location.

When an agent retrieves items from memory, search starts from the current pointer location and spreads out in a spherical fashion. See figure 6. AWM doesn't include separate short-term and long-term memory stores. Short-term memory is the part of long term memory that is currently active. The model demonstrates recency and salience effects often associated with the concept of short-term memory since a belief recently encountered will be closer to the current memory pointer location. It also incorporates frequency effects since one is much more likely to retrieve an item that is stored in multiple memory locations[Lan75]. AWM provides a natural account of a decay in the salience of a proposition stored in memory,

since as the memory pointer moves, propositions get displaced, and thus are no longer salient.




Figure 6: Searching Memory: Spreading Search for a Fixed Radius

The Design-World simulation environment allows us to vary the radius of the search sphere (attentional capacity) and examine how this affects performance (score) on the task. As shown in figure 7, scores increase as attentional capacity (search radius) increases. As capacity increases, the benefits of increased capacity tail off and eventually the agents achieve optimal performance. This means that whenever a certain performance is desired, it can be guaranteed, given enough resources. However if action is required at low levels of resources, some reasonable behavior is produced.

## 4 Discourse Strategies

An agent’s communication strategy determines which discourse actions an agent has available, and in particular whether they produce IRUs in certain discourse situations. The hypotheses about the benefits of IRUs are tested by comparing the performance of

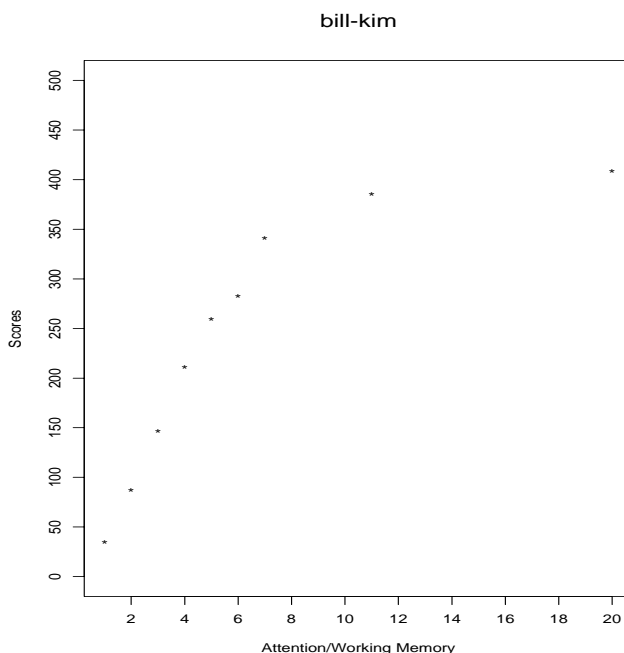


Figure 7: Baseline, Two Attention Implicit Agents

agents that systematically produce IRUs in certain discourse situations with the performance of agents who don’t.

This paper reports on a comparison of results of simulations of three strategies that exemplify the relations between Attention and Inference. In the first strategy, Baseline, the agents do not produce IRUs. The task performance results for the baseline strategy is shown in figure 7. An excerpt of a dialogue where both agents use an implicit attention strategy is shown in example 3. Agents communicate in an artificial language, and a gloss from this language is produced by adhoc generation strategies, and shown in *italics*. Each corresponding utterance in the artificial language is shown just after the gloss.<sup>3</sup>

(3) BILL: *Then, let’s put the fuchsia couch in the study.*

1:(propose agent-bill agent-kim option-1552: put-act (agent-bill

<sup>3</sup>These excerpts are all from experimental runs in which AWM was set to 8 for both agents.

fuchsia couch room-1) )

KIM: *No, instead let's put the green lamp in the study.*

2:(reject agent-kim agent-bill option-1553: put-act (agent-kim green lamp room-1))

BILL: *Then, let's put the green rug in the study.*

3:(propose agent-bill agent-kim option-1568: put-act (agent-bill green rug room-1))

KIM: *Then, let's put the purple chair in the study.*

4:(propose agent-kim agent-bill option-1589: put-act (agent-kim purple chair room-1))

Note here that opening, closing and acceptance are all done implicitly, and no extra information is provided along with a proposal that can be used in evaluating the proposal. Agents must base their reasoning on what has been said and what they can retrieve from AWM.

In the two comparison strategies, agents' communication strategies include IRUs whose content consists of facts that the speaker intends the hearer to use in reasoning. These strategies are meant to abstractly model the discourse functions identified by the distributional analysis presented in [Wal92], and exemplified by examples 1 and 2.

In the first strategy, OPEN-BEST, an agent tells another agent that he has a piece of furniture, even though the other agent already knows that. An excerpt of a dialogue where both agents are using the OPEN-BEST Attention strategy is shown in 4, where in 4-6, agent OP22 produces an IRU, shown in CAPS. Other utterances are shown in *italics*.

(4) OP22: AGENT-KIM HAS PURPLE CHAIR

6:(say agent-op22 agent-op2 bel-1997: has (agent-kim purple chair) )

OP22: *Then, let's put the purple couch in the study*

7:(propose agent-op22 agent-op2 option-1418: put-act (agent-kim purple couch room-1))

OP2: SO, WE'VE AGREED TO PUT THE PURPLE COUCH IN THE STUDY.

8:(close agent-op2 agent-op22 intended-1396: put-act (agent-kim purple couch room-1)) . . .

This strategy is a rough approximation of introducing facts that an agent is using in reasoning into the context, even when they are IRUs. The content of the IRU is determined by OP22's means-end reasoning, in which using the purple chair was considered because the purple chair is the second highest scoring piece of those that OP22 was reasoning about. OP22 then proposed to use his highest scoring piece.<sup>4</sup>

A second strategy relating Attention and Inference is the CONSEQUENCE-WARRANT strategy in which an agent makes explicit the reasons why the current proposal is a good one, at the time that the proposal is made. An excerpt from a dialogue in which agent IEI uses this strategy is given in 5:

(5) IEI: PUTTING IN THE GREEN RUG IS WORTH 56.

1:(say agent-ie1 agent-ie2 bel-2379: score (option-1694: put-act (agent-bill green rug room-1) 56))

IEI: *Then, let's put the green rug in the study.*

2:(propose agent-ie1 agent-ie2 option-1694: put-act (agent-bill green rug room-1))

---

<sup>4</sup>A variation on the OPEN-BEST strategy, called the OPEN-ALL strategy consists of an agent telling another agent about ALL the pieces that she considered during means-end reasoning before making her proposal. This strategy was also shown to have beneficial effects at low AWM settings.



IEI2: PUTTING IN THE GREEN LAMP IS WORTH 55.

3:(say agent-ie12 agent-ie1 bel-2404: score (option-1712: put-act (agent-kim green lamp room-1) 55))

IEI2: *Then, let's put the green lamp in the study.*

4:(propose agent-ie12 agent-ie1 option-1712: put-act (agent-kim green lamp room-1))

In both strategies, the IRU makes a proposition salient which was already mutually believed, but not available for reasoning in the current segment. The performance of agents using these strategies will be discussed in section 5.

## 5 Results

As shown in figure 8, the OPEN-BEST Attention strategy improves performance on the task. These IRUs consist of additional utterances at the beginning of a dialogue segment in which one agent reminds the other agent that he has another piece that is also of high value (the next highest valued piece), in addition to the piece that he is proposing to use to contribute to the current intention. At memory values of 1, 2, 3 and 4, the OPEN-BEST strategy improves performance. At higher memory values of 5 and 6, the two strategies are beginning to converge and have completely converged when AWM is 11.

The temporarily lower scores before convergence may show that too much redundancy can be detrimental with attention-limited agents. What appears to be happening in these cases is that the IRUs displace facts from memory that the agent could have used to perform better on the design task. At AWM of 11, both agents are performing optimally because they can do more extensive search.

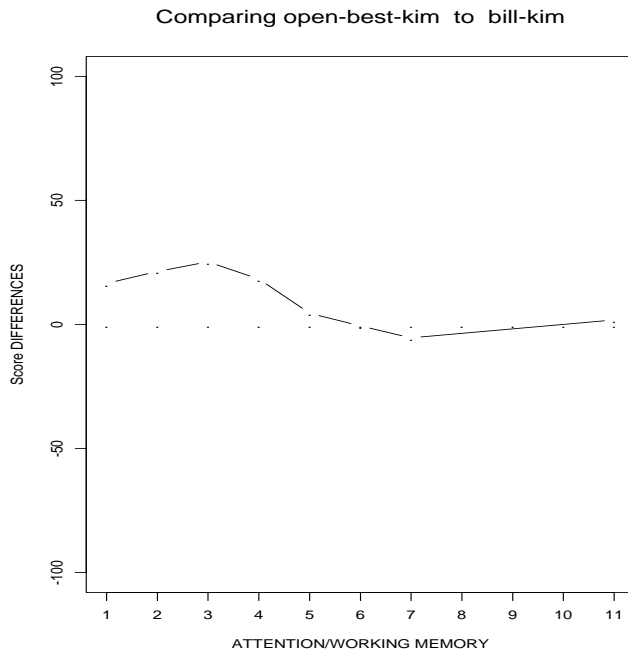


Figure 8: Differences of Baseline compared with OPEN-BEST Attention Strategy

As shown in figure 9, the CONSEQUENCE-WARRANT strategy increases performance for AWM settings above 4. Performance continues to increase as AWM settings increase, because agents need to do much less searching of memory for facts to be used in reasoning when this dialogue strategy is used. There is little benefit however at the lowest settings for AWM, presumably because the propositions in question are already reasonably available or because IRUs may displace other facts from memory that would be more useful.

## 6 Conclusions

This paper describes a simulation environment, Design-World, developed in order to explore the interaction of agents' discourse strategies and their resource limitations. This environment includes a psychologically motivated account of agents' limited attentional capacity. This paper reports results on the effects of limited attention on agents' inferential capacity and concomitant scores on a




Figure 9: Differences of Baseline compared with the Consequence-Warrant Strategy

collaborative task. I demonstrate that discourse strategies that include IRUs can increase performance on the task when agents are resource-limited.

The strategies discussed here are used consistently by an agent who doesn't reason about what to do in each discourse situation. In other words, the agents here apply the same strategy in all situations and with all conversational partners. A more fine-grained and better targeted use of IRUs would presumably demonstrate more beneficial effects, but this depends on adaptive strategies, and must be left to future work.

Future work also includes exploration of other discourse strategies and their effect on inference with other variations in the simulation environment such as task complexity, inferential complexity, agents with different capabilities, and uncertainty in communication.

## References

- [BIP88] Michael Bratman, David Israel, and Martha Pollack. Plans and resource bounded practical reasoning. *Computational Intelligence*, 4:349–355, 1988.
- [Car92] Jean C. Carletta. *Risk Taking and Recovery in Task-Oriented Dialogue*. PhD thesis, Edinburgh University, 1992.
- [CGRJ92] Alison Cawsey, Julia Galliers, Steven Reece, and Karen Sparck Jones. Automating the librarian: A fundamental approach using belief revision. Technical Report 243, Cambridge Computer Laboratory, 1992.
- [CL91] Phil Cohen and Hector Levesque. Teamwork. *Nous*, 25:11–24, 1991.
- [Coh78] Phillip R. Cohen. On knowing what to say: Planning speech acts. Technical Report 118, University of Toronto; Department of Computer Science, 1978.
- [Gaz79] Gerald Gazdar. *Pragmatics : implicature, presupposition and logical form*. Academic Press, 1979.
- [GS90] Barbara J. Grosz and Candace L. Sidner. Plans for discourse. In *Cohen, Morgan and Pollack, eds. Intentions in Communication*, MIT Press, 1990.
- [Jos78] Aravind K. Joshi. Some extensions of a system for inference on partial information. In *Pattern Directed Inference Systems*, pages 241–257. Academic Press, 1978.
- [JW81] Aravind K. Joshi and Scott Weinstein. Control of inference: Role of some aspects of discourse structure - centering. In *Proc. International*

- Joint Conference on Artificial Intelligence*, pages 385–387, 1981.
- [JWW84] Aravind K. Joshi, Bonnie Lynn Webber, and Ralph M. Weischedel. Preventing false inferences. In *COLING84: Proc. 10th International Conference on Computational Linguistics.*, pages 134–138, 1984.
- [KST82] Daniel Kahneman, Paul Slovic, and Amos Tversky. *Judgment under uncertainty : heuristics and biases*. Cambridge University Press, 1982.
- [Lan75] Thomas K. Landauer. Memory without organization: Properties of a model with random storage and undirected retrieval. *Cognitive Psychology*, pages 495–531, 1975.
- [Lan87] Amy Lansky. Gemplan. *Computational Intelligence*, page 1, 1987.
- [Mil56] G. A. Miller. The perception of speech. In M. Halle, editor, *For Roman Jakobson*, pages 353–359. The Hague, Mouton Publishers, 1956.
- [PHW82] Martha Pollack, Julia Hirschberg, and Bonnie Webber. User participation in the reasoning process of expert systems. In *AAAI82*, 1982.
- [Pol90] Martha E. Pollack. Plans as complex mental attitudes. In *Cohen, Morgan and Pollack, eds. Intentions in Communication*, MIT Press, 1990.
- [PR90] Martha E. Pollack and Marc Ringuette. Introducing the tile-world: Experimentally evaluating agent architectures. In *AAAI90*, pages 183–189, 1990.
- [Sid92] Candace Sidner. A collaborative model of negotiation dialogues. *AAAI Workshop on Heterogeneous Agents*, 1992.
- [Sol92] Miriam Solomon. Scientific rationality and human reasoning. *Philosophy of Science*, September 1992.
- [Sta78] Robert C. Stalnaker. Assertion. In Peter Cole, editor, *Syntax and Semantics, Volume 9: Pragmatics*, pages 315–332. Academic Press, 1978.
- [Wal92] Marilyn A. Walker. A model of redundant information in dialogue: the role of resource bounds. Technical Report IRCS-92-25, University of Pennsylvania Institute for Research in Cognitive Science, 1992.
- [WGR93] Steve Whittaker, Erik Geelhoed, and Elizabeth Robinson. Shared workspaces: How do they work and when are they useful? *To Appear in IJMMMS*, 1993.
- [WJ82] Bonnie Lynn Webber and Aravind Joshi. Taking the initiative in natural language database interaction: Justifying why. Department of Computer and Information Science, University of Pennsylvania, 1982.
- [WS88] Steve Whittaker and Phil Stenton. Cues and control in expert client dialogues. In *Proc. 26th Annual Meeting of the ACL*, pages 123–130, 1988.
- [WW90] Marilyn A. Walker and Steve Whittaker. Mixed initiative in dialogue: An investigation into discourse segmentation. In *Proc. 28th Annual Meeting of the ACL*, pages 70–79, 1990.