The Institute For Research In Cognitive Science

**Informational Redundancy and
Resource Bounds in Dialogue**

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

P

E

N

N

Informational Redundancy and Resource Bounds in Dialogue

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

Abstract

The view of language as action is well-entrenched in computational theories of discourse, however previous work has not paid particular attention to the fact that agents' resource bounds must affect the way that dialogue proceeds just as it affects agents' behavior in general. This paper presents results on the role of discourse strategies in improving performance on collaborative tasks when agents are resource-limited. The resource limitation addressed here is limits in inferential capacity, and the interaction of discourse strategies with this limitation is explored through a dialogue simulation environment, Design-World, in which inferential capacity can be parameterized. The results reported here demonstrate that discourse strategies that incorporate redundant information can improve the performance of inference-limited agents to the point of logical omniscience.

1 Introduction

The view of language as action is well-entrenched in computational theories of discourse, however previous work has not paid particular attention to the fact that agents' resource bounds must affect the way that dialogue proceeds just as it affects agents' behavior in general. This paper focuses on the relation of agents' resource bounds to a particular discourse phenomenon, INFORMATIONALLY REDUNDANT utterances (IRUs), defined below:

> An utterance u_i is INFORMATIONALLY REDUNDANT in a discourse situation $\mathcal{S}$ if u_i expresses a proposition p_i, and p_i is already entailed in $\mathcal{S}$.

INFORMATIONAL REDUNDANCY is a paradox according to the language as action view; it appears that agents are performing actions whose effects have already been achieved. Yet IRUs account for approximately 12% of the utterances in a large corpus of naturally occurring multi-agent problem-solving dialogues[22, 18]. Examples of IRUs and some results from a distributional analysis of the corpus will be discussed further in section 1.1. I argue elsewhere that IRUs are not **communicatively** redundant, and serve three communicative functions[21]:

- Attitude: to address the assumptions underlying the inference of mutual understanding and acceptance

1

- Consequence: to augment the evidence supporting mutual beliefs that certain inferences are licensed, or to provide a context in which an inference is more easily made

- Attention: to manipulate the locus of attention of the discourse participants by making or keeping a proposition salient

Furthermore, each of these communicative functions is related to a particular way in which agents are resource-bounded.[1] Attitude follows from agents' autonomy and the resultant potential for costly misunderstandings and disagreements[9]. Consequence follows from the fact that agents are not logically omniscient[14, 13, 16]. Attention follows from the fact that agents have limited attentional capacity[1]. This paper focuses on the interaction of Consequence and Attention.

In order to explore the relationship of resource-bounds and dialogue strategies that incorporate IRUs, I have developed a computational simulation environment called Design-World. The Design-World environment and task is presented in section 2. Like Tileworld, Design-World is structured around the IRMA agent architecture for resource-bounded agents[4, 19]. The Design-World simulation replaces action in the environment, with communicative acts that propose future actions. The IRMA architecture supports the separation of bounds on means-end reasoning and deliberation, which is necessary to test the hypotheses about the communicative functions of IRUs. Limited means-end reasoning is modeled by discrete Design-World parameter settings constraining how many inferences are made.

Design World adds a model of limited attention to the belief module of the IRMA architecture. Section 3 describes this parameterizable model of attention that simulates both recency and frequency effects[15]. The major goal of the experiments is to test whether dialogue strategies that use IRUs are more efficient when agents have limited attention and limited means-end reasoning. The dialogue strategies will be discussed in section 5. Sections 5.1.1 and 5.2.1 describe the results of experiments to explore the interactions of IRUs and resource-bounds.

1.1 Examples of IRUs

A distributional analysis was performed on a corpus of joint problem-solving dialogues from a radio talk show for financial advice[18].[2] The utterance(s) that originally added the propositional content of the IRU to the discourse situation is called the IRU's ANTECEDENT. In the examples below, IRUs are marked with CAPS whereas their antecedents are given in *italics*.

An example of a Consequence IRU is given in 1.

(1)   (15) h. Oh no. *I R A's were available as long as you are not a participant in an existing pension*
         (16) j. Oh I see. Well I did work, *I do work for a company that has a pension*
         (17) h. ahh. THEN YOU'RE NOT ELIGIBLE FOR EIGHTY ONE

In example 1, (17) is inferrable via logical inference from (15) and (16) and yet the talk show host Harry (h) makes this inference explicit. Consequence IRUs are related to the fact that agents are not logically omniscient. If agents make inferences explicit, this demonstrates that they made a particular inference,

---

[1] Individual IRUs may simultaneously address one or more of these functions: if inferences are focused on what is currently attended to, as seems likely[14], then limits on inferencing follow directly from limits on attention.

[2] This consisted of a qualitative analysis of 471 IRUs from 54 dialogues, and a quantitative analysis of 176 IRUs from 24 dialogues.

and also ensures that another agent can draw that inference.

Consequence IRUs also include cases where agents restate premises that are already mutually believed so that inferences dependent on these premises are easily made. Consider example 2:

(2)   (42) h. Now what is your income situation?
       (43) m. *We're both retired* and our income for the year is about um 24.. about 26 thousand
       (44) h. Have you other securities than stock? Have you any bonds or certificates?
       (45) m. Yes yes we do- we have some certificates oh about uh 15 - 20 thousand, not much, we're not rich - and we have a house completely paid for, have some land in the Poconos, completely paid for - and uh that's actually the extent of our uh
       (46) h. Ok - on the proceeds of that GM stock
       (47) m. Yes
       (48) h. I'd like to see you put that into two different southern utilities
       .
       (clarification of southern utilities)
       .
       (58) h. And those I think that you will find that that will give you a good return,
       YOU ARE RETIRED
       and that is primarily what you are looking for, in addition I think that eventually those stocks will rise as well
       (59) m. uh huh

Beginning with (46) Harry (h) suggests a course of action. In (58) he paraphrases to Mary (m) a fact that she told him in (43), *you are retired*. However Harry's statement in this context leads to the inference that the fact that Mary and her husband are retired is

a WARRANT for the intended act, i.e. the suggested course of action is motivated by that fact.

Example 3 demonstrates the communicative function of Attention. IRUs are often used to set the context and ensure that agents are jointly attending to the same concepts[23]. Example 3-(22) consists of two clauses that both realize propositions discussed earlier in 3-(4) to (6).

(3)   (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
       (23) e. yes
       (24) h. and I don't like putting all my eggs in one basket.....

These Attention IRUs redirect the listener's attention to the previously discussed propositions. A hypothesis tested in the simulation is that dialogues where one conversant produces Attention IRUs, which function to recreate the previous context, are more cost effective than dialogues where agents are always made

to retrieve those propositions from memory. Note that the initial phrase of (22), *but as far as the certificates are concerned*, should be enough to direct the listener's attention to that context, so that the IRUs in (22) could conceivably have been retrieved from memory based on the initial phrase alone.

	SALIENT	DISPLACED
Repetitions	50	11
Paraphrases	32	28
Inferences	24	8

Figure 1: The distribution of IRUs with respect to the discourse status of their antecedent

A distributional analysis of the location of IRUs with respect to the location of their antecedents is given in the table shown in figure 3. The results show that logical inferences usually have salient antecedents (24 out of 32). This fact provides weak support for the hypothesis that limited inference is mainly determined by limited attention. This analysis also shows that paraphrases are more likely than the other types to have displaced antecedents, (28 out of 60). These paraphrases are often examples similar to 2 where a fact discussed earlier is used as a WARRANT for a course of action discussed later. The other most common cases of remote paraphrases are those that manipulate attention as in example 3. These two kinds of cases are clearly related: in both cases a fact discussed earlier in the dialogue is used later in reasoning.

2 Design World Simulation Experiments

The Design-World simulation environment consists of the floor plan for a house[24]. The Design-World task consists of two agents who must carry out a dialogue in order to come to an agreement on a design for a floor plan with two rooms. Each agent has different pieces of furniture and there are more pieces available than needed to accomplish the task. Furniture items are of 5 types: couch, table, chair, lamp and rug. Each furniture item has a color and point value. A design for a room consists of any four pieces from these types. The score associated with a furniture item supports the calculation of utility of including that item in the design plan. The agents attempt to maximize the score achieved together by their design. Figure 2 shows a potential initial state for a dialogue.

Performance is measured by the score associated with the final design as compared with the costs to achieve this score. 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. This makes it possible to explore in a principled manner the trade-offs associated with different dialogue strategies.

	ROOM # 1		ROOM # 2

AGENT A's PIECES:

RED COUCH	30 pts.
BLUE TABLE	25 pts.
YELLOW COUCH	15pts.
BLUE CHAIR	15 pts.
RED RUG	20 pts
YELLOW LAMP	10 pts.

AGENT B's PIECES:

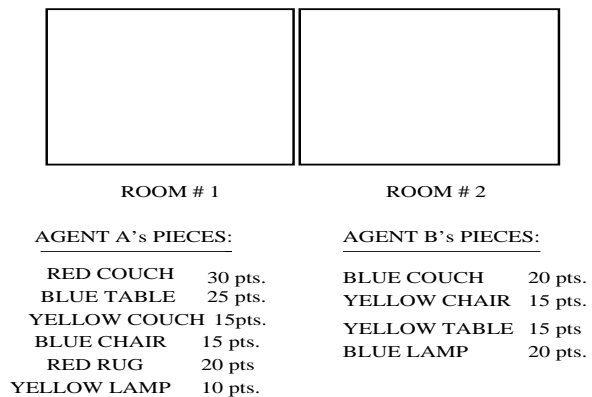| BLUE COUCH | 20 pts. |
YELLOW CHAIR	15 pts.
YELLOW TABLE	15 pts
BLUE LAMP	20 pts.

Figure 2: Potential Initial State for the Design-World Task

In Design-World, unlike Tileworld, the environment consists of both the domain-based task and the other agent. In the Design-World task, the potential intended acts, ie. options, are domain level intended acts such as putting a particular furniture piece into a particular

4

room. These options may eventually correspond to proposals that an agent makes to another agent, depending on the outcome of the deliberation process. Messages from other agents create options for the receiver that they didn't previously know about, which are then subject to deliberation.

2.1 Domain and Plan Structure

Each agent starts with private beliefs as to what pieces of furniture she has and what colors these pieces are. Agents share beliefs about which pieces of furniture exist and how many points they are worth. These are represented with domain state predicates. The only domain actions are:

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

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

Goals are to design a room and to match the color of the furniture in a room whenever possible since a MATCHED-PAIR doubles the points of the pieces involved. Agents are from a set of two, Ann and Bob, as well as Ann and Bob as a pair, A&B. The ?Time variable for each action and state is either NOW or LATER.

In the Design-World simulation, the agents share beliefs about what it means to have a COLLABORATIVE PLAN to achieve a goal(See also [23, 7, 12]):

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

1. MB A&B (Intend A $\vee$ B ($\alpha_1 \wedge \ldots \alpha_n$))

2. $\forall \alpha_i$ (MB A&B (Contribute α_i (Achieve A&B Goal))

3. MB A&B (Max-Utility ($\alpha_1 \wedge \ldots \alpha_n$) (Achieve A&B Goal))

The use of the Max-Utility constraint in the formulation of collaborative planning incorpo-

rates the role of deliberation in planning future actions[8]. 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 make 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, or acceptance/rejection is postponed by ASKING for more information. A schema for agents' messages is given in figure 3.[3]
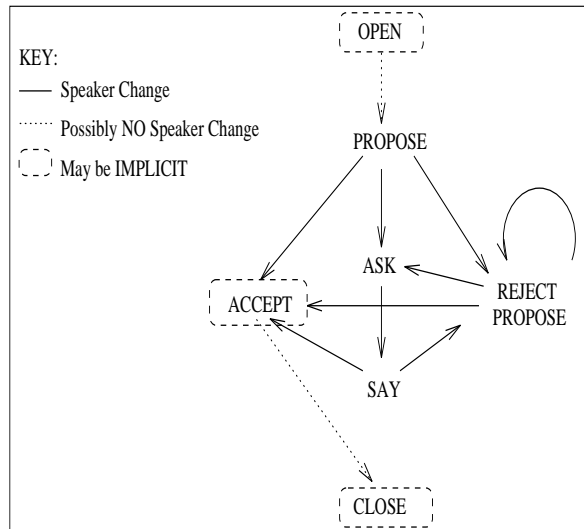

Figure 3: Dialogue Actions for the Design-World Task:Version 1

The dialogues that agents engage in in Design-World are MIXED-INITIATIVE[23]. Both agents reason about the domain and both agents have goals so either agent can make a PROPOSAL. Furthermore, even though agents are working together to perform the task, there is no

---

[3]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. See [5, 20] for alternate utterance taxonomies used in simulations. The hypothesis is that the effects of resource-bounds will operate independently of the range of utterance level intentions and other strategies available to an agent.

guarantee that they will accept one another's proposals. Each agent reasons about whether to accept or reject a proposal based on her view of the proposed action's utility and thus there is no enforced collaboration [2].[4] Acceptance generates the mutual beliefs necessary for the agents to have a COLLABORATIVE PLAN, but acceptance can be done implicitly or explicitly.[5]

As shown in figure 3, a rejection always includes a counter-proposal. This simplification is due to the fact that the decision to reject a proposal is based on comparing it with other options an agent knows about, and another option must be of greater utility to support a rejection. Although it is possible for an agent to respond to a rejection with another rejection, all conflicts are eventually resolved since agents share a global utility function based on the score associated with a particular intention. This is illustrated in the excerpt from a sample dialogue in 4 below:[6]

(4)   BILL: Let's put the green couch in the study
KIM: No, instead let's put the green lamp in the study
BILL: Putting in the green couch is worth 54
BILL: No, instead let's put the green

couch in the study
KIM: Putting in the green lamp is worth 55
KIM: No, instead let's put the green lamp in the study
BILL: Then, let's put the green couch in the study

DESIGN WORLD COLLABORATIVE PLAN : 280 points

Figure 4: Potential Final State for Design-World Task: Represents the Collaborative Plan Achieved by the Dialogue, assuming Matched-Pair goals

The communicative strategies associated with making and accepting or rejecting proposals and with acquiring and providing information on which deliberation is based is discussed in section 5. Figure 4 shows a potential final collaborative plan for the task as a result of the dialogue, based on the initial state given in figure 2.

3   Limited Memory/Attention Model

There is an abundant amount of evidence that attention/working memory (AWM) capacity is limited in human agents[17, 1, 3]. However, it is not at all clear what exactly these limits are and how they interact with agents' reasoning capabilities.

The AWM model used in Design World is a

---

[4]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, B may assume that A's proposal isn't worth much or B may ask A for the value of A's proposal. However, if B has no options available to satisfy the same goal, B can accept A's proposal without knowing its exact value.

[5]Utterances that indicate acceptance explicitly are most often Attitude IRUs. The account of Attitude IRUs is that they address the assumptions underlying the inference of mutual belief, which depends on the evidence that the addressee provides as to the effect of an utterance[25, 23, 6]. When the underlying assumptions are supported by explicit evidence in the dialogue, the mutual beliefs that are dependent on them are less defeasible(see also [11, 10]). This weak model of mutual belief is is discussed elsewhere[22, 21].

[6]Agents communicate with a propositional language. This dialogue gloss is generated automatically using adhoc procedures from the propositional language that agents communicate with.

very simple model that incorporates both recency and frequency effects, and which has been shown to account for a broad range of empirical results on human memory and learning[15]. This model consists of a three dimensional space in which beliefs are stored in chronological sequence according to the current location of a moving memory pointer. The next memory pointer location is never more than two steps away from the current location. However, of the distance 2 locations, the location to be used next is randomly selected. Thus beliefs acquired near the same time will tend be closer together, and beliefs stored in multiple locations have a greater chance of being retrieved.


Figure 5: Baseline-ALL: Two ALL Inference, Consequence Implicit, Attention Implicit Agents, Averages of 100 runs

When items are to be retrieved from memory, search starts from the current memory pointer location and spreads out in a spherical fashion. AWM is modeled as the radius of this search sphere. As shown in figure 5, AWM alone provides a main effect on scores in the simulations. The points plotted are stable averages of 100 runs at each parameter setting. In addition, when AWM is set to 11, the agents achieve optimal performance. This means that whenever a level of performance is desired, it can be guaranteed, but if action is required at low levels of resources, some reasonable behavior is produced.

The AWM/Score plot shown in figure 5 is a Baseline for ALL inference agents (Baseline-ALL) against which to compare Explicit Consequence and Explicit Attention strategies discussed below.

4   Limiting Consequence

Another parameter of the Design World simulation is a resource-bound on the number of inferences that agents may make. This is realized with 3 discrete settings: (1) agents store NO inferences in memory; (2) agents store HALF of the possible inferences; (3) agents store ALL of the possible inferences.
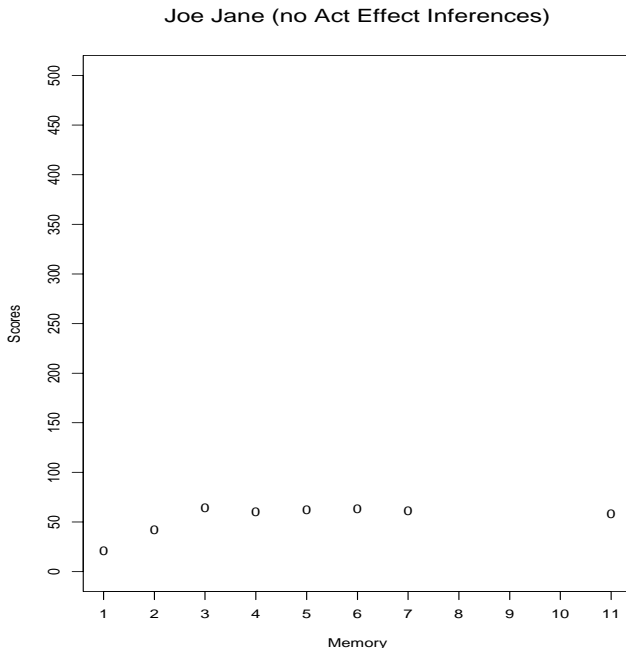

Figure 6: Baseline-NO: Two NO Inference, Consequence Implicit and Attention Implicit Agents, Averages of 100 runs

Figure 5 plots performance as a function of

AWM when both agents are ALL inference agents, i.e. logically omniscient. Figure 6 shows the scores achieved by two agents who store NO inferences. The AWM/Score plot shown in figure 6 is a Baseline for NO inference agents (Baseline-NO) against which the Explicit Consequence and Explicit Attention strategies will be compared. The low scores of these NO inference agents are one result of the large number of invalid steps in their plans because they don't infer that once they have used a piece of furniture in a plan step, they no longer have it.

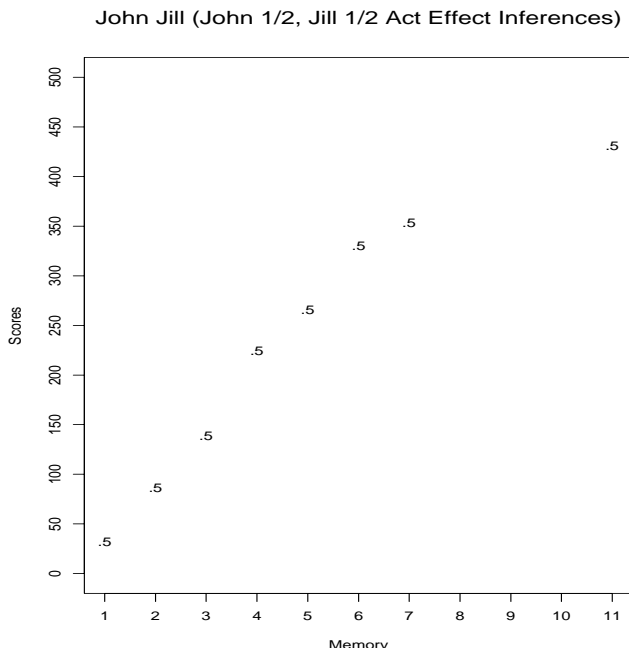John Jill (John 1/2, Jill 1/2 Act Effect Inferences)

Figure 7: Two HALF Inference Agents, Consequence and Attention Implicit

Figure 7 shows the scores achieved by two agents who each make HALF of the possible inferences. Surprisingly, these two HALF inference agents do not perform significantly worse than the Baseline-ALL agents shown in figure 5.

The next section introduces the communication strategies that interact with these basic resource bounds. Section 5.1.1 and 5.2.1 will describe the results of experiments exploring these interactions.

5 Communication 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 agents that systematically produce IRUs in certain discourse situations as compared with agents that don't. The communicative strategies associated with the Consequence and Attention IRU classes are discussed below.

5.1 Attention Strategies

There are two Explicit Attention strategies: Open-All and Open-Best. These are to be compared with an agent using the Baseline strategy shown in figure 5 who makes no utterances that explicitly indicate attentional state. Both the Open-All and Open-Best strategies explicitly indicate the opening and closing of a segment of the dialogue.[7] Basic Open and Close statements for the simulation are given below:

(Say A B
(Open (Achieve A&B (Design Room-1))))

(Say A B
(Close (Achieve A&B (Design Room-1))))

The Open-All strategy consists of the Basic Open plus additional statements in which the agent communicates every single fact he thinks of, as he reasons about the current intention. This is shown in example 5. The IRUs are given in CAPS:

(5)   .
        .
        Bill: Now, let's put something in the living room.
        Bill: BILL HAS THE GREEN RUG

---

[7]Segments are task-defined as the utterances needed for reasoning and discussion about one intention.

Bill: BILL HAS THE GREEN
COUCH
Bill: BILL HAS THE RED COUCH
Bill: BILL HAS THE RED LAMP
Bill: BILL HAS THE GREEN
CHAIR
Bill: BILL HAS THE RED CHAIR
Bill: BILL HAS THE RED TABLE
Bill: Let's put the green rug in the
study.
.
.

The Open-Best strategy consists of the Open statement above, plus two additional statements in which the agent communicates the facts used in reasoning about the two best options she has identified. See example 6:

(6)   Kim: Now, let's put something in the study
Kim: KIM HAS THE BLUE CHAIR
Kim: KIM HAS THE BROWN TABLE
Kim: Then, let's put the brown table in the study

This strategy is common in the human-human dialogues for this type of task, where the agents concentrate on using the highest scoring pieces first[24]. The following section discusses the results of using these strategies.

5.1.1   Results for Attention

The AWM/Score plots shown in figures 8 and 9 are plots of the **differences** between the Open-Best and the Open-All strategies and the Baseline strategy shown in figure 5. At AWM values of 1 to 4, Open-Best is a better strategy than Baseline. However there is an AWM/Strategy interaction and with AWM at 5, Baseline does better. The hypothesized explanation for this fact is that IRUs can displace facts from memory that might have contributed more to performance.
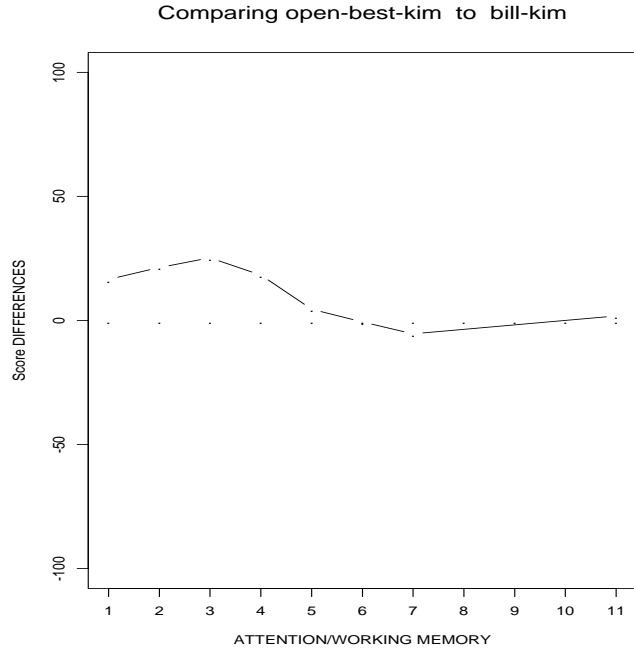

Figure 8: Two ALL Inference Agents, One is Open-Best

As shown in figure 9, the Open-All strategy also does better than the Baseline strategy.
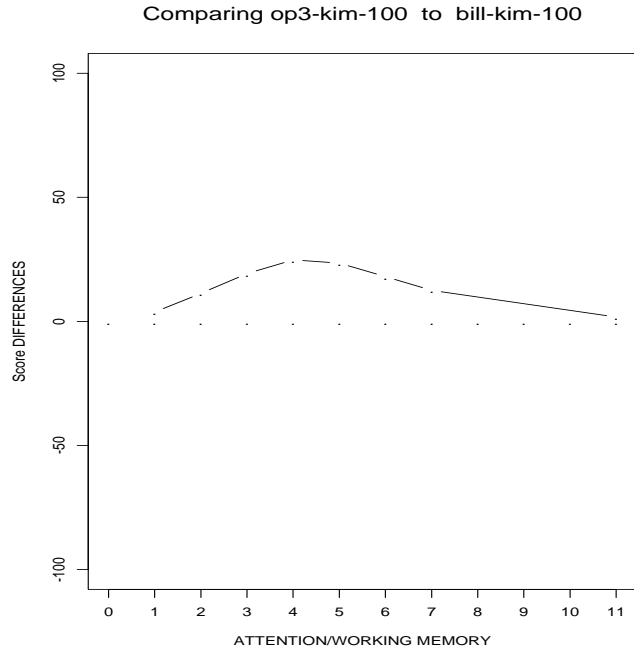

Figure 9: Two ALL Inference Agents, One is Open-All

9

However, there is no significant difference between the Open-All and the Open-Best strategies.

5.2 Consequence Strategies

A second set of simulations tests two different Consequence strategies against the Baseline. Remember that the Baseline is an implicit Consequence strategy, i.e. no IRUs related to Consequence are communicated; the AWM/Score plot for this strategy is shown in figure 5. The other two strategies are explicit Consequence strategies: (1) Consequence-Inference (CI): An agent makes inferences explicit that follow from facts already mutually believed; (2) Consequence-Warrant (CW): An agent may say facts that are already mutually believed when these facts provide a WAR-RANT for actions currently under consideration. The CI strategy parallels the naturally occurring example in 1, whereas the CW strategy parallels the naturally occurring example in 2.

In the simulation the CI strategy is realized by including consequences of actions in CLOSE statements:

(Close A B
(Put A&B Red-Couch Room-1 Later))
(Say A B (NOT (HAS A Red-Couch))))

The CW strategy is realized by including facts to be used in reasoning about acceptance in PROPOSE statements, e.g. the Score information in the proposal below might serve as a WARRANT for intending the action described in the proposal.

(Propose A B
(Put A&B Red-Couch Room-1 Later))
(Say A B (Score Red-Couch 30)))

In both cases, the utterance includes additional information that is already mutually believed since in the CI strategy an agent might have inferred that if a piece is used the agent no longer has it, and in the CW strategy both agents know the point values associated with furniture items at the beginning of the task.

5.2.1 Results for Consequence

First, CI was tested with inference ALL agents. A hypothesis was that even though inference ALL agents would definitely have inferred the additional fact included in the CI strategy, increasing the frequency of this fact in memory could increase robustness and potentially increase scores. However, this doesn't seem to be the case.

Next, the CI strategy was tested on NO inference agents. The effect of only one agent in the dialogue employing the CI strategy is shown by the **difference** plot in figure 10. This figure shows the difference between the score of two agents, one of whom employs the Consequence-Inference strategy, and the NO inference agents shown in figure 6. In this situation, as expected, the strategy of making inferences explicit improves scores significantly.
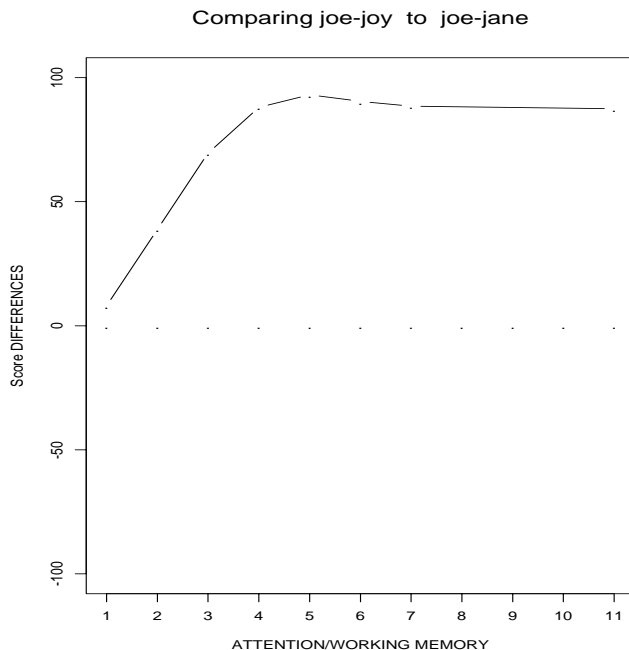

Figure 10: Two NO Inference Agents, Joy is Consequence-Inference

If both agents are HALF Inference agents, then the CI strategy has no effect. Figure 11 shows the difference plot of this situation with that shown in figure 7. This might be due to the fact that there is no significant difference between the performance of HALF inference agents and ALL inference agents.
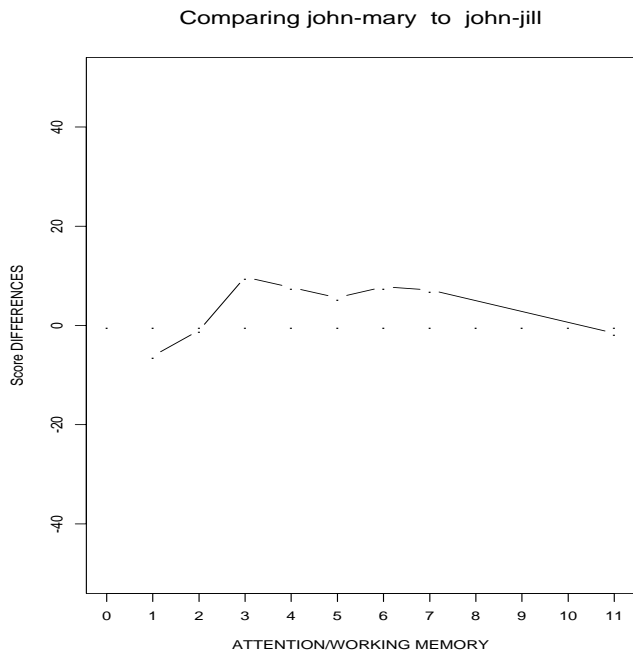
Comparing john-mary to john-jill

Figure 11: Two HALF Inference Agents, Mary is Consequence Inference

Furthermore, if both agents employ the CI strategy, as shown in figure 12, they perform as well as logically omniscient agents. Figure 12 plots the differences in scores of these agents with those shown in figure 5. This result is of interest since it demonstrates that situations exist in which NO inference agents perform as well as (or better than) logically omniscient agents. Thus situations in which NO inference agents are strongly preferred will be those when it is possible for a large number of inferences to be drawn from a fact, of which only one is intended or relevant. In this case, the CI strategy in combination with NO inference agents is vastly more efficient.

Next consider the effects of the CW strategy. Surprisingly, the CW strategy has no effect

**Comparing joy2-joy to bill-kim**

Figure 12: Differences between two NO Inference Agents when both are Consequence-Inference, and two Logically omniscient agents who never produce IRUs

on scores. A prediction might have been that at low AWM values, agents would make bad decisions about which pieces to use because they could not retrieve the information about the scores of pieces. Nevertheless, it appears that agents can perform just as well either not knowing the scores of pieces or retrieving them from their own beliefs. Figure 13 plots the number of retrievals associated with the CW strategy (w) as compared with the Baseline.

However, as shown in figure 13, a closer inspection of the costs involved with achieving these scores shows that agents using the CW strategy are doing much less processing in terms of retrievals from memory than Baseline agents.[8] An appropriate composite score measure should be derivable in which this difference in processing costs is reflected in scores.

---

[8]I used a log scale for retrievals from memory because the underlying distribution wasn't normal.
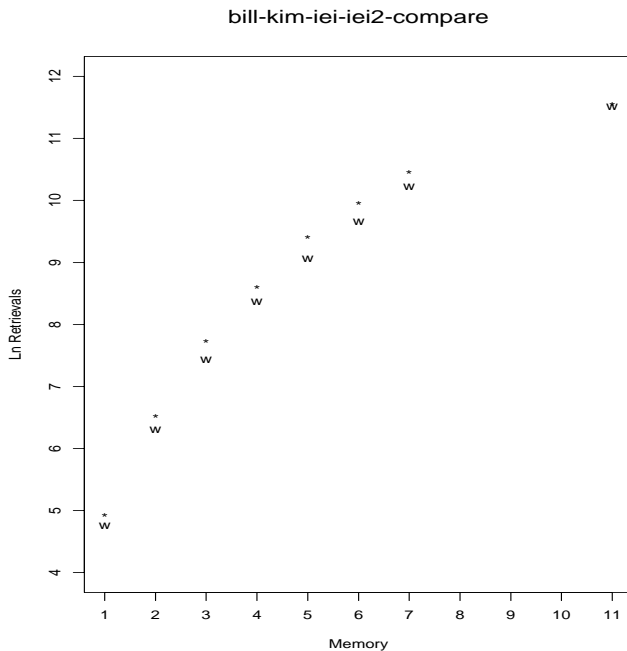
Figure 13: LN Retrievals of two ALL Inference Agents, * = Implicit, W=Consequence-Warrant

6 Conclusions

The analysis of the function of IRUs in dialogue highlights the fact that the information-exchange view of dialogue is too narrow. I proposed that IRUs can function to address limitations in agents' processing capabilities and provided support for this claim through a number of simulation experiments. This paper presents an operationalization of resource-bounded agents along different resource dimensions and tested the effects of different classes of IRUs in overcoming different classes of resource bounds. Future work includes increasing the complexity of the inferences involved in performing the DesignWorld task and developing procedures for evolving communication strategies in tandem with other variables in the communicative situation such as costs for communication, inference and retrieval and the effects of uncertainty in communication.

References

[1] J. R. Anderson and G. H. Bower. *Human Associative Memory*. V.H. Winston and Sons, 1973.

[2] Robert Axelrod. *The Evolution of Cooperation*. Basic Books, 1984.

[3] Alan Baddeley. *Working Memory*. Oxford University Press, 1986.

[4] Michael Bratman, David Israel, and Martha Pollack. Plans and resource bounded practical reasoning. *Computational Intelligence*, 4:349–355, 1988.

[5] Jean C. Carletta. *Risk Taking and Recovery in Task-Oriented Dialogue*. PhD thesis, Edinburgh University, 1992.

[6] Herbert H. Clark and Edward F. Schaefer. Contributing to discourse. *Cognitive Science*, 13:259–294, 1989.

[7] Phil Cohen and Hector Levesque. Teamwork. *Nous*, 25:11–24, 1991.

[8] Jon Doyle. Rationality and its roles in reasoning. *Computational Intelligence*, November 1992.

[9] Julia R. Galliers. A theoretical framework for computer models of cooperative dialogue, acknowledging multi agent conflict. Technical Report 172, University of Cambridge, Computer Laboratory, 1989.

[10] Julia R. Galliers. Autonomous belief revision and communication. In P. Gardenfors, editor, *Belief Revision*, pages 220 – 246. Cambridge University Press, 1991.

[11] Julia R. Galliers. Cooperative interaction as strategic belief revision. In M.S. Deen, editor, *Cooperating Knowledge Based Systems*, pages 148 – 163. Springer Verlag, 1991.

[12] Barbara J. Grosz and Candace L. Sidner. Plans for discourse. In *Cohen, Morgan and Pollack, eds. Intentions in Communication, MIT Press*, 1990.

[13] Gilbert Harman. *Change of View*. MIT PRESS, 1986.

[14] Daniel Kahneman, Paul Slovic, and Amos Tversky. *Judgment under uncertainty : heuristics and biases*. Cambridge University Press, 1982.

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

[16] Hector Levesque. A logic of implicit and explicit belief. In *AAAI84*, pages 198–202, 1984.

[17] G. A. Miller. The perception of speech. In M. Halle, editor, *For Roman Jakobsen*, pages 353–359. The Hague, Mouton Publishers, 1956.

[18] Martha Pollack, Julia Hirschberg, and Bonnie Webber. User participation in the reasoning process of expert systems. In *AAAI82*, 1982.

[19] Martha E. Pollack and Marc Ringuette. Introducing the Tileworld: Experimentally Evaluating Agent Architectures. In *AAAI90*, pages 183–189, 1990.

[20] Candace Sidner. Using discourse to negotiate in collaborative activity: An artificial language. *AAAI Workshop on Cooperation among Heterogeneous Agents*, 1992.

[21] 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. Dissertation Proposal.

[22] Marilyn A. Walker. Redundancy in collaborative dialogue. In *Fourteenth International Conference on Computational Linguistics*, 1992.

[23] 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.

[24] Steve Whittaker, Erik Geelhoed, and Elizabeth Robinson. Shared workspaces: How do they work and when are they useful? *To Appear in IJMMS*, 1993.

[25] Steve Whittaker and Phil Stenton. Cues and control in expert client dialogues. In *Proc. 26th Annual Meeting of the ACL, Association of Computational Linguistics*, pages 123–130, 1988.