Design-World:
A testbed of communicative action and resource limits

Marilyn A. Walker
Mitsubishi Electric Research Labs
201 Broadway
Cambridge, Ma. 02139
walker@merl.com

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

1 Introduction

Design-World is a simulation testbed developed in order to support experiments on the relationship of communicative action to agents' processing limitations. This testbed can be used as a tool for teaching both technical and methodological concepts in AI.

On the technical side, running Design-World experiments introduces the student to concepts of dialogue and how it is structured, natural language generation and planning communicative acts, multi-agent collaboration as a team during problem-solving, and the effect of resource limits on agent action.

On the methodology side, Design-World introduces concepts in experimental AI, such as forming and testing a hypothesis via simulation experiments, methods for evaluating statistical significance of simulations run with different parameter settings, and the role and limitations of simulation testbeds in AI theory development (Pollack and Ringuette, 1990; Hanks, Pollack, and Cohen, 1993).

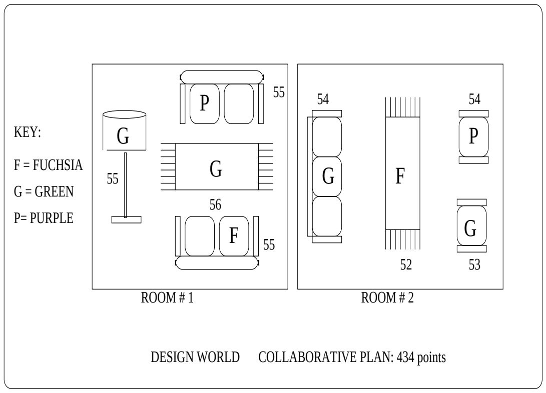

Figure 1: Potential Final State for Design-World Task

Design-World is based on a simple two-agent collaborative task to negotiate an agreement on the furniture layout of a two room house. Two simulated, highly parameterizable agents converse until a final plan is negotiated. The model of agent interaction is based on the concept of a **team** of agents doing a task together (Joshi, Webber, and Weischedel, 1986; Cohen and Levesque, 1991; Grosz and Sidner, 1990), who retain their autonomy throughout the task, and thus must negotiate each plan step (Galliers, 1989; Walker and Whittaker, 1990; Chu-Carrol and Carberry, 1994; Sidner, 1994).

Figure 1 shows a potential final plan negotiated via a dialogue. A partial dialogue generated from running Design-World is given below, showing both the artificial language that the agents communicate with as well as a gloss generated from that language in italics:

(1) 1: BILL: *First, let's put the green rug in the study.*
(propose agent-bill agent-kim option-10: put-act (agent-bill green rug room-1))

2: KIM: *Then, why don't we put the green lamp in the study?*
(propose agent-kim agent-bill option-33: put-act (agent-kim green lamp room-1))

3: BILL: *Then, let's put the green couch in the study.*
(propose agent-bill agent-kim option-45: put-act (agent-bill green couch room-1))

4: KIM: *No, instead let's put in the purple couch.*
(reject agent-kim agent-bill option-56: put-act (agent-kim purple couch room-1))

The simulation environment uses only the basic features of CommonLisp and is completely portable. Dialogues can be simulated interactively or noninteractively. Functions are provided to facilitate running multiple simulations noninteractively with performance statistics written to a file for comparisons with other simulations. In addition, a graphical interface to Design-World has been implemented using the Garnet Toolkit from the CMU software archive, to provide visual output of two artificial agents' progress on the task. A planned extension is to allow the experimenter to interact graphically as an agent doing the task.

The remainder of this paper describes the Design-World environment and the simulation parameters, and suggests potential student assignments using Design-World as a teaching tool.

2 Design-World

The purpose of the Design-World dialogue simulation testbed is to support experiments on the interaction between distinct parameters in multi-agent communication for problem solving: (1) available processing resources, (2) communicative strategy, and (3) the version of the task. The main processing resource that is varied is the amount of working memory that agents have access to, which indirectly affects their inferential ability (Norman and Bobrow, 1975; Walker, 1993). Agent communicative strategies vary as to how explicit the agents are, how much information they include at various points in the dialogue, or how much they leave to the other agent to retrieve or infer. The task parameters vary how fault intolerant the task is, how many inferences are required to perform the task, and the degree to which agents must coordinate on the beliefs underlying the final negotiated plan for the task.

The agent architecture for deliberation and means-end reasoning is based on the IRMA architecture, also used in the TileWorld simulation environment (Bratman, Israel, and Pollack, 1988; Pollack and Ringuette, 1990), with the addition of

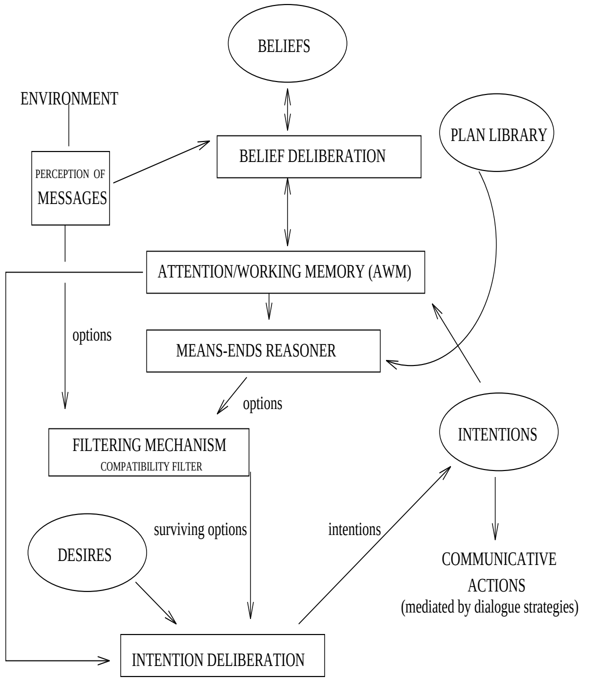

Figure 2: Design-World version of the IRMA Agent Architecture for Resource-Bounded Agents with Limited Attention (AWM)

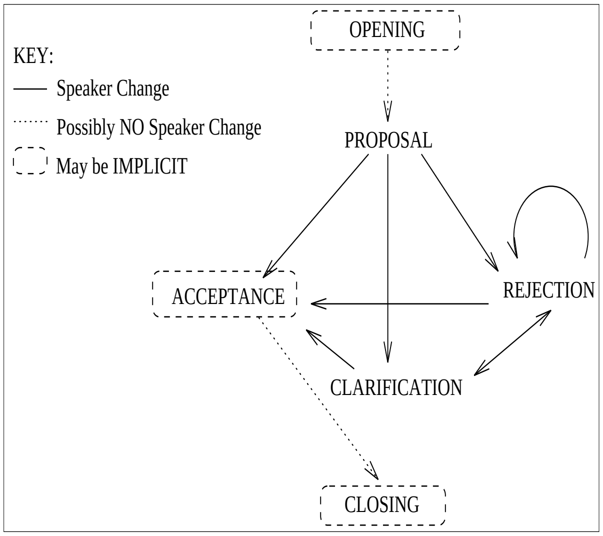

Figure 3: Discourse Actions for the Design-World Task

a model of limited Attention/Working memory, AWM (Landauer, 1975; Baddeley, 1986). See figure 2.

At the beginning of a Design-World simulation, both agents know about a set of furniture items. These furniture items provide the basis for means-end reasoning about how to do the task. Each furniture item has a type, a color and a point value. The point value provides the basis for agent deliberation and the color property is the source of inferences and additional task constraints.

In order to negotiate a plan, the agents alternate turns in which they (1) process incoming messages whose content are beliefs and (potential) intentions; (2) means-end reason about options that can satisfy the current intention; (3) deliberate according to potential domain options identified and the current dialogue situation (e.g., whether a proposal has just been made), based on beliefs about the proposed action's utility (Doyle, 1992); (4) respond according to the deliberation result and their communicative strategy.

2.1 Communicative Strategy Parameters

One of the major issues in communication planning is how agents decide what to say and when to say it, and what information they must access in order to do so. In Design-World, agents are parametrized by communicative strategy parameters that determine what the agent says whenever it is in a particular discourse situation. By varying other parameters we can determine when each strategy is beneficial.

Communicative strategies are built out of the primitive communicative acts of SAY, ASK, PROPOSE, ACCEPT, REJECT, OPEN, CLOSE, similar to those used in other work (Whittaker and Stenton, 1988; Carletta, 1992; Sidner, 1994; Chu-Carrol and Carberry, 1994). Communicative acts are composed into higher level discourse acts of OPENINGS, PROPOSALS, REJECTIONS, ACCEPTANCES or CLOSINGS as shown in figure 3.

Communicative strategies are defined according to the plan expansions the agent has for each type of discourse act (Moore and Paris, 1993; Young, Moore, and Pollack, 1994). Each plan step must be realized communicatively by a propose communicative act, but each plan-step may also generate additional discourse acts as shown in figure 3, depending on the agent's communicative strategy. Currently implemented strategies include:

- All Implicit: any discourse act (e.g. opening, acceptance and closing) that can be inferred by the other agent is left implicit. In addition proposals contain the minimum information possible.
- Close Consequence: explicitly close the discussion of a proposal by making an explicit closing statement, e.g. *Okay, we're done with that.*, and then adding an extra assertion that makes the inferences explicit that follow from the acceptance of the proposal, e.g. *That gets us 130 points so far.*
- Explicit Warrant: Each proposal expands to a propose communicative act and a statement that provides a reason to accept a proposal, e.g. *The green chair is worth 50 points* (Webber and Joshi, 1982). In the experiments presented below both agents knew the warrant information at the onset of the dialogue, but due to resource limitations may not have had that information immediately accessible.
- Matched-Pair-Inference-Explicit: help the other agent make matched pair inferences by expanding the proposal discourse act to two communicative acts: a statement about an already intended act that can be used as a premise for inferring a match, followed by the proposed match.

These strategies are all based on strategies observed in human-human dialogues. Dialogue 1 illustrated the All-Implicit strategy. Implementing a new communicative strategy is simple: a new plan expansion for a discourse act is placed in the agents' plan library, and given a name. Then agents can be parameterized to use that strategy.

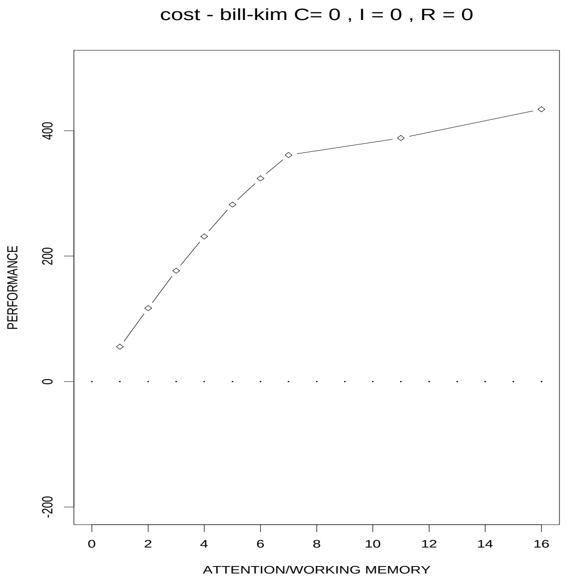

Figure 4: Effect of AWM parameterization on two All-Implicit Agents

2.2 Resource Parameters

One type of parametrization of available resources involves setting the size of AWM (attention/working memory) at a value between 1 and 16, where 1 is a severely limited agent and agents with AWM of 16 can access all of their memory.

Another type of resource parameterization is linked to the architecture for resource-bounded agents shown in Figure 2. The model for the simulation is that the cost of the processes involved with doing the task should be minimized over the team of agents, as in Clark's model of LEAST COLLABORATIVE EFFORT (Clark and Schaefer, 1989). Each module in the architecture contributes to the total effort. Thus the simulation tracks how many inferences, how many memory retrieval operations, and how many messages were required to complete the task for **both** agents.

Since the effort for each of these processes varies according to the implementation of the agent architecture, agents' retrieval, inference and communicative costs are parameterized by (1) COMMCOST: cost of sending a message; (2) INFCOST: cost of inference; and (3) RETCOST: cost of retrieval from memory. Collaborative effort is then defined as:

COLLABORATIVE EFFORT =
(COMMCOST × total messages for both agents)
+ (INFCOST × total inferences for both agents)
+ (RETCOST × total retrievals for both agents)

Performance on the task is: Task Defined RAW SCORE – COLLABORATIVE EFFORT.

2.3 Task Parameters

The raw score, as just defined, varies according to the task definition:

- Standard: The raw score for a collaborative plan is simply the sum of the furniture pieces in the plan with the values of the invalid steps subtracted out.
- Zero-Invalid: Give a zero score to a collaborative plan with any invalid steps in it, reflecting a binary division between task situations where the substeps of the task are independent from those where they are not.
- Zero-Nonmatching-Beliefs: Give a zero score to a collaborative plan when the agents disagree on the reasons for having adopted a particular intention. These reasons are the WARRANTS for the intention.
- Matched-Pair-Tasks: Agents must match the color of furniture items within or between rooms in order to get points for them. The final score is the scores for the matched-pair optional goals that both agents inferred.

The Zero-Invalid task is a fault intolerant version of the Standard task, whereas the other three variations all test various aspects of interagent coordination with different communicative strategies. Zero-Nonmatching beliefs penalizes agents who do not agree on the reasons underlying agreed upon goals, whereas the two Matched-Pair tasks penalize agents who do not coordinate on inferences.

3 Running Experiments and Evaluating Results

Each time the simulation is run, it provides a data point for how well a team of agents with particular communicative strategies and a particular amount of available resources did on a particular task.

Design-World includes tools for exploring the effects of different parameter settings. The software includes S routines for plotting performance visually for how a pair of agents perform over a range of resource settings. For example, figure 4 shows the effect of parametrizing AWM on two All-Implicit Agents over the range of 1 . . . 16, when all resources are free.

The software also includes S routines for plotting performance **differences** between two sets of agents doing the task, and for evaluating the statistical significance of these differences, over the full range of resource (AWM) settings. To illustrate the evaluation of performance differences, all the plots below will compare the performance of two Explicit-Warrant agents with two All-Implicit agents. Figure 5 plots performance differences between these two agent types over all resource settings when there is no cost associated with accessing memory (retcost = 0). An implementation of the Kolmogorov-Smirnov (KS) two sample test allows the experimenter to evaluate the statistical significance of performance differences for different simulation runs (Siegel, 1956). The KS two-sample test is a test of whether two independent samples have been drawn from the same population or from populations with the same distribution. For an experimental parameter to affect performance, it must make the cumulative performance distributions appear to have been drawn from different populations.

When comparing two strategies, a strategy is BENEFICIAL if the difference in distributions using the KS two-sample test is significant at $p < .05$, in the positive direction, for two or more AWM settings. A strategy is DETRIMENTAL if the differences go in the negative direction. For example, figures 5 and 6 show the effect on performance differences of two sets of agents when the cost of retrieval is varied. Figure 6 shows that there is no difference in the agents performance at limited AWM of 3,4,5 but that when resources have some cost and agents can potentially use more resources (have higher AWM), then communicative strategies that help them

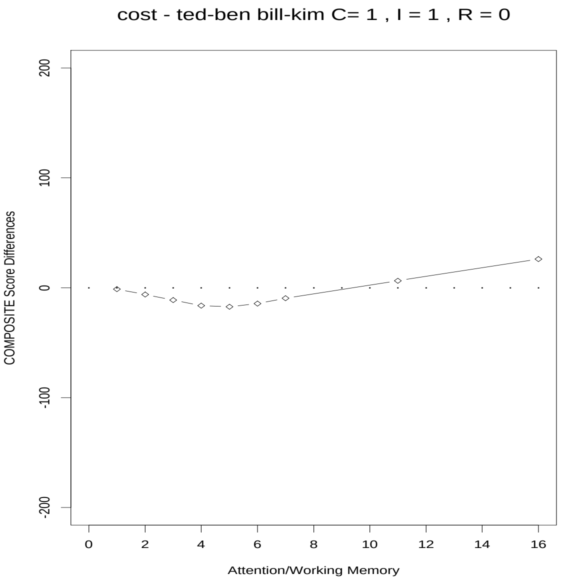

Figure 5: If Retrieval is Free, Explicit-Warrant is detrimental at AWM of 3,4,5: Strategy 1 of Two Explicit-Warrant agents and strategy 2 of two All-Implicit agents: Task = Standard, commcost = 1, infcost = 1, retcost = 0

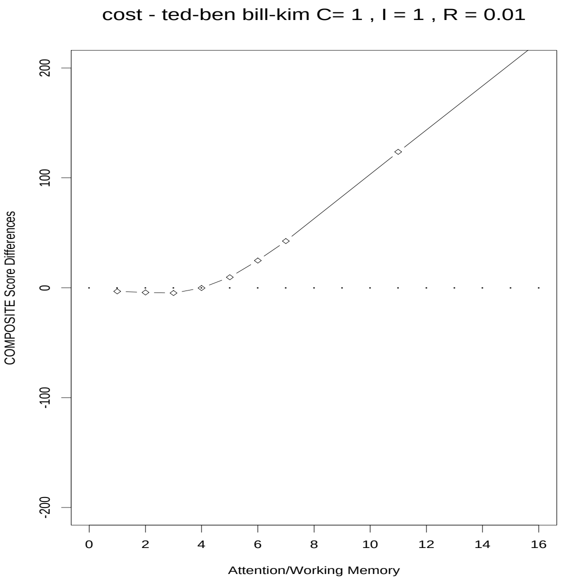

Figure 6: If Retrieval is Not Free, Explicit-Warrant is beneficial at higher AWM. Strategy 1 is two Explicit-Warrant agents and strategy 2 is two All-Implicit agents: Task = Standard, commcost = 1, infcost = 1, retcost = .01

limit resource use can be beneficial (Walker, 1994a). Figure 7 shows that a high communication cost makes Explicit-Warrant a detrimental strategy for the Standard task.

Figure 8 shows the effect of varying the task parameter. Under the Zero-Nonmatching-Beliefs, the explicit-warrant strategy is highly beneficial since that task requires agents to be coordinated on the reasons underlying their intentions. In this task, this strategy is beneficial even when communication costs dominate all other processing costs, in contrast with figure 7 for the Standard task (Walker, 1994b).

4 Using Design World for Teaching

Potential assignments using Design-World could include:

- Students set up agents with particular parameter settings, select a version of the task and then run experiments of the agents doing the task. After running the experiments, the student evaluates how well the different types of agents do on the task in comparison with one another by using the visualization methods and by running statistical tests.
- Students add a new communicative strategy to an agent's repertoire, e.g. always accept the other agent's proposal or always ask a question rather than disagree. Then students must evaluate how agents with this strategy perform on the outcome measures for the task, and test how the performance varies depending on other factors, such as the information distribution between the agents (Walker and Whittaker, 1990).
- Students add a decision algorithm as an extension of a communication strategy. The algorithm bases the decision to communicate a proposition on the total cognitive effort and students evaluate how agents using this decision strategy compare to two extreme strategies where in one case the decision is always "yes" and in the other it is always "no" (Jordan and Walker, 1995).

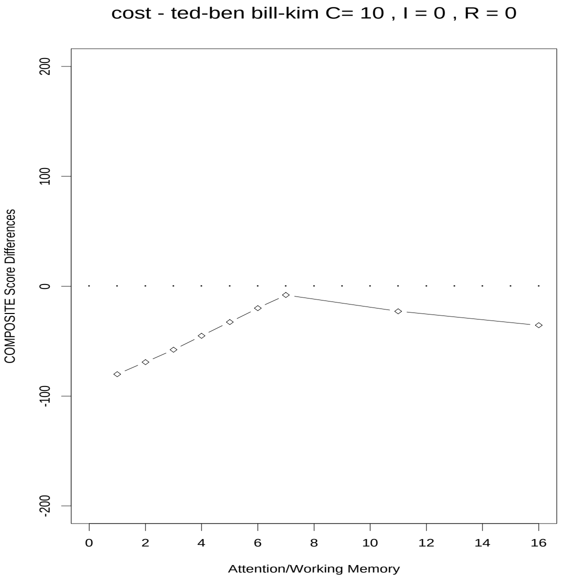

Figure 7: If Communication is expensive, Explicit-Warrant is detrimental. Strategy 1 is two Explicit-Warrant agents and strategy 2 is two All-Implicit agents: Task = Standard, commcost = 10, infcost = 0, retcost = 0

- Students change one of the modules of the simulation and compare simulation runs with the original versus the new model in order to determine the effect of the model on the simulation results. Modules that would be easy to change are the working memory model, the belief mediation component, the domain planner, the deliberator.

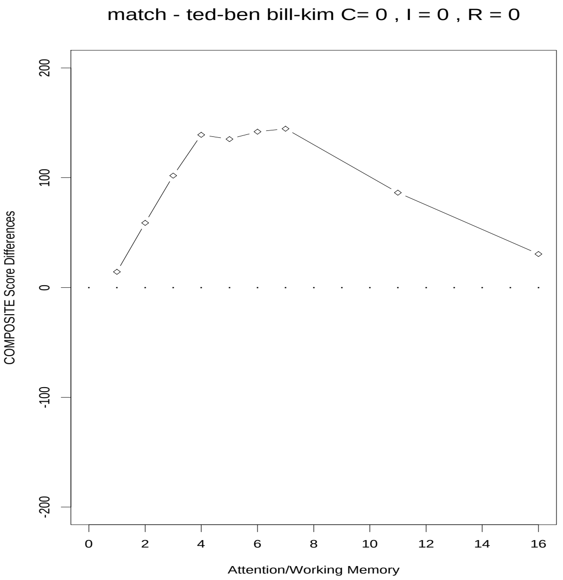

Figure 8: Explicit-Warrant is beneficial for Zero-Nonmatching-Beliefs task. Strategy 1 is two Explicit-Warrant agents and strategy 2 is two All-Implicit agents: commcost = 0, infcost = 0, retcost = 0

- Students give the deliberator access to different evaluation metrics for plans, and add negotiation strategies so that agents can resolve conflicts when applying different evaluators to plans.
- Students add a new parameter to the simulation, such as a new task definition.

Bibliography

Alan Baddeley. 1986. *Working Memory.* Oxford University Press.

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

Jean C. Carletta. 1992. *Risk Taking and Recovery in Task-Oriented Dialogue.* Ph.D. thesis, Edinburgh University.

Jennifer Chu-Carrol and Sandra Carberry. 1994. A plan-based model for response generation in collaborative task-oriented dialogue. In *AAAI 94*, pages 799–805.

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

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

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

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

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

Steve Hanks, Martha Pollack, and Paul Cohen. 1993. Benchmarks, testbeds, controlled experimentation and the design of agent architectures. *AI Magazine*, December.

Pamela W. Jordan and Marilyn A. Walker. 1995. Designing resource sensitive agents for communication. Technical Report 1995 Series, Mitsubishi Electric Research Labs.

Aravind K. Joshi, Bonnie Webber, and Ralph M. Weischedel. 1986. Some aspects of default reasoning in interactive discourse. Technical Report MS-CIS-86-27, University of Pennsylvania.

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

Johanna D. Moore and Cécile L. Paris. 1993. Planning text for advisory dialogues: Capturing intentional and rhetorical information. *Computational Linguistics*, 19(4).

Donald A. Norman and Daniel G. Bobrow. 1975. On data-limited and resource-limited processes. *Cognitive Psychology*, 7(1):44–6.

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

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

Sidney Siegel. 1956. *Nonparametric Statistics for the Behavioral Sciences.* McGraw Hill.

Marilyn A. Walker and Steve Whittaker. 1990. Mixed initiative in dialogue: An investigation into discourse segmentation. In *Proc. 28th Annual Meeting of the ACL*, pages 70–79.

Marilyn A. Walker. 1993. *Informational Redundancy and Resource Bounds in Dialogue.* Ph.D. thesis, University of Pennsylvania.

Marilyn A. Walker. 1994a. Discourse and deliberation: Testing a collaborative strategy. In *COLING 94*.

Marilyn A. Walker. 1994b. Experimentally evaluating communicative strategies: The effect of the task. In *AAAI94*.

Bonnie Webber and Aravind Joshi. 1982. Taking the initiative in natural language database interaction: Justifying why. In *Coling 82*, pages 413–419.

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

R. M. Young, J. D. Moore, and M. E. Pollack. 1994. Towards a principled representation for discourse plans. In *Proc 16th Annual Meeting of the Cognitive Science Society*, pages 946–951.