

Redundancy in Collaborative Dialogue

Appeared in Coling92

Marilyn A. Walker

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

Abstract

In dialogues in which both agents are autonomous, each agent deliberates whether to accept or reject the contributions of the current speaker. A speaker cannot simply assume that a proposal or an assertion will be accepted. However, an examination of a corpus of naturally-occurring problem-solving dialogues shows that agents often do not explicitly indicate acceptance or rejection. Rather the speaker must infer whether the hearer understands and accepts the current contribution based on indirect evidence provided by the hearer's next dialogue contribution. In this paper, I propose a model of the role of informationally redundant utterances in providing evidence to support inferences about mutual understanding and acceptance. The model (1) requires a theory of mutual belief that supports mutual beliefs of various strengths; (2) explains the function of a class of informationally redundant utterances that cannot be explained by other accounts; and (3) contributes to a theory of dialogue by showing how mutual beliefs can be inferred in the absence of the master-slave assumption.

1 Introduction

It seems a perfectly valid rule of conversation not to tell people what they already know. Indeed, Grice's QUANTITY maxim has often been interpreted this way: *Do not make your contribution more informative than is required*[5]. Stalnaker, as well, suggests that *to assert something that is already presupposed is to attempt to do something that is already done*[14]. Thus, the notion of what is *informative* is judged against a background of what is *presupposed*, i.e. propositions that all conversants assume are mutually known or believed. These propositions are known as the COMMON GROUND[10, 5].

The various formulations of this 'no redundancy' rule

permeate many computational analyses of natural language and notions of cooperativity. However consider the following excerpt from the middle of an advisory dialogue between Harry (h), a talk show host, and Ray (r) his caller¹.

Example 1:

- (6) r. uh 2 tax questions.
one: since April 81 we have had an
85 year old mother living with us.
her only income has been social security
plus approximately \$3000 from a
certificate of deposit and i wonder
what's the situation as far as
claiming her as a dependent or does
that income from the certificate of
deposit rule her out as a dependent?
(7) h. yes it does.
(8) r. IT DOES.
(9) h. YUP THAT KNOCKS HER OUT.
.....

In standard information theoretic terms, both (8) and (9) are REDUNDANT. Harry's assertion in (9) simply paraphrases what was said in (7) and (8) and so it cannot be adding beliefs to the common ground². Furthermore, the truth of (9) cannot be in question, for instead of (9), Harry could not say *Yup, but that doesn't knock her out*. So why does Ray (r) in (8) REPEAT Harry's (h) assertion of *it does*, and why does Harry PARAPHRASE himself and Ray in (9)?

My claim is that informationally redundant utterances (IRU's) have two main discourse functions: (1) to provide EVIDENCE to support the assumptions underlying the inference of mutual beliefs, (2) to CENTER a proposition, ie. make or keep a proposition salient[6]. This paper will focus on (1) leaving (2) for future work.

¹These examples come from the talk show for financial advice, *Speaking of Your Money*, on WCAU in Philadelphia. This corpus was collected and transcribed by Martha Pollack and Julia Hirschberg[12].

²(8) is not realized with a rising question intonation. This will be discussed in section 6.1.

*This research was partially funded by ARO grant DAAL03-89-C0031PRI and DARPA grant N00014-90-J-1863 at the University of Pennsylvania, by Hewlett Packard, U.K., and by an NSF award for 1991 Summer Institute in Japan

First consider the notion of evidence. One reason why agents need EVIDENCE for beliefs is that they only have partial information about: (1) the state of world; (2) the effects of actions; (3) other agent's beliefs, preferences and goals. This is especially true when it comes to modelling the effects of linguistic actions. Linguistic actions are different than physical actions. An agent's prior beliefs, preferences and goals cannot be ascertained by direct inspection. This means that it is difficult for the speaker to verify when an action has achieved its expected result, and so giving and receiving evidence is critical and the process of establishing mutual beliefs is carefully monitored by the conversants.

The characterization of IRU's as informationally redundant follows from an axiomatization of action in dialogue that I will call the DETERMINISTIC MODEL. This model consists of a number of simplifying assumptions such as: (1) Propositions are either believed or not believed, (2) Propositions representing beliefs and intentions get added to the context by the unilateral action of one conversant, (3) Agents are logically omniscient. (4) The context of a discourse is an undifferentiated set of propositions with no specific relations between them. I claim that these assumptions must be dropped in order to explain the function of IRU's in dialogue.

Section 2 discusses assumption (1); section 3 shows how assumption (2) can be dropped; section 4 discusses (3); section 4.1 shows that some IRU's facilitate the inference of relations between adjacent propositions.

2 Mutual Beliefs in a Shared Environment

The account proposed here of how the COMMON GROUND is augmented, is based on Lewis's SHARED ENVIRONMENT model for common knowledge[10, 2]. In this model, mutual beliefs depend on evidence, openly available to the conversants, plus a number of underlying assumptions.

Shared Environment Mutual Belief Induction Schema

It is mutually believed in a population P that Ψ if and only if some situation \mathcal{S} holds such that:

1. Everyone in P has reason to believe that \mathcal{S} holds.
2. \mathcal{S} indicates to everyone in P that everyone in P has reason to believe that \mathcal{S} holds.
3. \mathcal{S} indicates to everyone in P that Ψ .

The situation \mathcal{S} , used above in the mutual belief induction schema, is the context of what has been said. This schema supports a weak model of mutual beliefs, that is more akin to mutual assumptions or mutual suppositions[13]. Mutual beliefs can be inferred based on some evidence, but these beliefs may depend on underlying assumptions that are easily defeasible. This model can be implemented using Gallier's theory of autonomous belief revision and the corresponding system[4].

A key part of this model is that some types of evidence provide better support for beliefs than other types. The types of evidence considered are categorized and ordered based on the source of the evidence: **hypothesis** < **default** < **inference** < **linguistic** < **physical**(See [2, 4]). This ordering reflects the **relative** defeasibility of different assumptions. Augmenting the strength of an assumption thus decreases its relative defeasibility.

A claim of this paper is that one role of IRU's is to ensure that these assumptions are supported by evidence, thus decreasing the defeasibility of the mutual beliefs that depend on them[4].

Thus mutual beliefs depend on a defeasible inference process. All inferences depend on the evidence to support them, and stronger evidence can defeat weaker evidence. So a mutual belief supported as an **inference** can get defeated by **linguistic** information. In addition, I adopt an assumption that a chain of reasoning is only as strong as its weakest link:

Weakest Link Assumption: The strength of a belief P depending on a set of underlying assumptions a_i, \dots, a_n is $\text{MIN}(\text{Strength}(a_i, \dots, a_n))$

This seems intuitively plausible and means that the strength of belief depends on the strength of underlying assumptions, and that for all inference rules that depend on multiple premises, the strength of an inferred belief is the weakest of the supporting beliefs.

This representation of mutual belief differs from the common representation in terms of an iterated conjunction[11] in that: (1) it relocates information from mental states to the environment in which utterances occur; (2) it allows one to represent the different kinds of evidence for mutual belief; (3) it controls reasoning when discrepancies in mutual beliefs are discovered since evidence and assumptions can be inspected; (4) it does not consist of an infinite list of statements.

3 Inference of Understanding

This section examines the assumption from the DETERMINISTIC MODEL that: (2) Propositions representing

beliefs and intentions get added to the context by the unilateral action of one conversant³. This assumption will also be examined in section 5.

The key claim of this section is that agents monitor the effects of their utterance actions and that the next action by the addressee is taken as evidence of the effect of the speaker's utterance⁴. That the utterance will have the intended effect is only a **hypothesis** at the point where the utterance has just been made, irrespective of the intentions of the speaker. This distinguishes this account from others that assume either that utterance actions always succeed or that they succeed unless the addressee previously believed otherwise[11, 8].

I adopt the assumption that the participants in a dialogue are trying to achieve some purpose[7]. Some aspects of the structure of dialogue arises from the structure of these purposes and their relation to one another. The minimal purpose of any dialogue is that an utterance be understood, and this goal is a prerequisite to achieving other goals in dialogue, such as commitment to future action. Thus achieving mutual belief of understanding is an instance of the type of activity that agents must perform as they collaborate to achieve the purposes of the dialogue. I claim that a model of the achievement of mutual belief of understanding can be extended to the achievement of other goals in dialogue.

Achieving understanding is not unproblematic, it is a process that must be managed, just as other goal achieving processes are[3]. Inference of mutual understanding relies upon some evidence, e.g. the utterance that is made, and a number of underlying assumptions. The assumptions are given with the inference rule below.

```
say(A, B, u, p) --A->
    understand(B, u, p) [evidence-type]
```

```
Assumptions =
{   copresent(A, B, u)      [evidence-type]
    attend(B, u)            [evidence-type]
    hear(B, u)              [evidence-type]
    bel(B, realize(u, p))    [evidence-type]
}
```

This schema means that when A says u to B intending to convey p, that this leads to the mutual belief that B understands u as p under certain assumptions. The assumptions are that A and B were copresent, that B was attending to the utterance event, that B heard the utterance, and that B believes that the utterance u realizes the intended meaning p.

The [evidence-type] annotation indicates the strength of evidence supporting the assumption. All of

³This is an utterance action version of the STRIPS assumption.

⁴Except for circumstances where it is clear that the flow of the conversation has been interrupted.

the assumptions start out supported by no evidence; their evidence type is therefore **hypothesis**. It isn't until **after** the addressee's next action that an assumption can have its strength modified.

The claim here is that one class of IRU's addresses these assumptions underlying the inference of mutual understanding. Each type of IRU, the assumption addressed and the evidence type provided is given in Figure 1. Examples are provided in sections 3.1 and 3.2.

It is also possible that A intends that BY saying u, which realizes p, B should make a certain inference q. Then B's understanding of u should include B making this inference. This adds an additional assumption:

```
bel(B, license(p,q)) [evidence-type]
```

Thus assuming that q was inferred relies on the assumption that B believes that p licenses q in the context.

Figure 1 says that prompts, repetitions, paraphrases and making inferences explicit all provide linguistic evidence of attention. All that prompts such as *uh huh* do is provide evidence of attention. However repetitions, paraphrases and making inferences explicit also demonstrate complete hearing. In addition, a paraphrase and making an inference explicit provides linguistic evidence of what proposition the paraphraser believes the previous utterance realizes. Explicit inferences additionally provide evidence of what inferences the inferrer believes the realized proposition licenses in this context.

In each case, the IRU addresses one or more assumptions that have to be made in order to infer that mutual understanding has actually been achieved. The assumption, rather than being a **hypothesis** or a **default**, get upgraded to a support type of **linguistic** as a result of the IRU. The fact that different IRU's address different assumptions leads to the perception that some IRU's are better evidence for understanding than others, e.g. a PARAPHRASE is stronger evidence of understanding than a REPEAT[3].

In addition, **any** next utterance by the addressee can upgrade the strength of the underlying assumptions to **default** (See Figure 1). Of course **default** evidence is weaker than **linguistic** evidence. The basis for these default inferences will be discussed in section 5.

3.1 Example of a Repetition

Consider example 1 in section 1. Ray, in (8), repeats Harry's assertion from (7). This upgrades the evidence for the assumptions of hearing and attention associated with utterance (7) from **hypothesis** to **linguistic**. The assumption about what proposition p7 is realized

Next Utterance	Assumption addressed	Evidence Type
PROMPT	attention	linguistic
REPEAT	hearing, attention	linguistic
PARAPHRASE	realize, hearing, attention	linguistic
INFERENCE	license, realize, hearing, attention	linguistic
IMPLICATURE	license, realize, hearing, attention	linguistic
ANY Next Utterance	copresence license, realize, hearing, attention	linguistic default

Figure 1: How the Addressee’s Following utterance upgrades the evidence underlying assumptions

by u7 remains a **default**. This instantiates the inference rule for understanding as follows:

```
say(harry, ray, u7, p7) --A->
    understand(Ray, u7, p7) [default]
```

```
Assumptions =
{   copresent(harry, ray, u7) [linguistic]
    attend(ray, u7)           [linguistic]
    hear(ray, u7)             [linguistic]
    bel(ray, realize(u7, p7)) [default]
}
```

Because of the **WEAKEST LINK** assumption, the belief about understanding is still a default.

3.2 Example of a Paraphrase

Consider the following excerpt:

Example 2:
(18) h. i see. are there any other children
beside your wife?
(19) d. no
(20) h. YOUR WIFE IS AN ONLY CHILD
(21) d. right. and uh wants to give
her some security

Harry’s utterance of (20) is said with a falling intonational contour and hence is unlikely to be a question. This utterance results in an instantiation of the inference rule as follows:

```
say(harry, ray, u20, p20) --A->
    understand(Ray, u20, p20) [linguistic]
```

```
Assumptions =
{   copresent(harry, ray, u7) [linguistic]
    attend(ray, u7)           [linguistic]
    hear(ray, u7)             [linguistic]
    bel(ray, realize(u7, p7)) [linguistic]
}
```

In this case, the belief about understanding is supported by **linguistic** evidence since all of the supporting assumptions are supported by linguistic evidence. Thus a paraphrase provides excellent evidence that an agent actually understood what another agent meant.

In addition, these IRU’s leave a proposition salient, where otherwise the discourse might have moved on to other topics. This is part of the **CENTERING** function of IRU’s and is left to future work.

4 Making Inferences Explicit

This section discusses assumption (3) of the deterministic model, namely that: Agents are logically omniscient. This assumption is challenged by a number of cases in naturally occurring dialogues where inferences that follow from what has been said are made explicit. I restrict the inferences that I discuss to those that are (a) based on information explicitly provided in the dialogue or, (b) licensed by applications of Gricean Maxims such as scalar implicature inferences[9].

For example the logical omniscience assumption would mean that if 1(a) and (b) below are in the context, then (c) will be as well since it is entailed from (a) and (b).

- (1) a. You can buy an I R A if and only if you do NOT have an existing pension plan.
 b. You have an existing pension plan.
 c. You cannot buy an I R A.

The following excerpt demonstrates this structure. Utterance (15) realizes 1a, utterance (16) realizes 1b, and utterance (17) makes the inference explicit that is given in 1c for the particular tax year of 1981.

Example 3:

- (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
 (18) j. i see, but i am for 82

After (16), since the propositional content of (17) is inferrable, the assumption that Harry has made this inference is supported by the **inference** evidence type:

`bel(H, license(p16, p17)) [inference]`

According to the model of achieving mutual understanding that was outlined in section 3, utterance (17) provides **linguistic** evidence that Harry (h) believes that the proposition realized by utterance (16) licenses the inference of (17) in this context.

`bel(H, license(p16, p17)) [linguistic]`

Furthermore, the context here consists of a discussion of two tax years *1981* and *1982*. Utterance (17) selects *eighty one*, with a narrow focus pitch accent. This implicates that there is some other tax year for which Joe is eligible, namely *1982*[9]. Joe's next utterance, *but I am for 82*, reinforces the implicature that Harry makes in (17), and upgrades the evidence underlying the assumption that (17) licenses (18) to **linguistic**.

4.1 Supporting Inferences

A subcase of ensuring that certain inferences get made involves the juxtaposition of two propositions. These cases challenge the assumption that: (4) The context of a discourse is an undifferentiated set of propositions with no specific relations between them. While this assumption is certainly not made in most discourse models, it is often made in semantic models of the context[14]. In the following segment, Jane (j) describes her financial situation to Harry (h) and a choice between a settlement and an annuity.

Example 4:

- (1) j. hello harry, my name is jane
 (2) h. welcome jane
 (3) j. i just retired december first,
 and in addition to my pension and
 social security, I have a
 supplemental annuity
 (4) h. yes
 (5) j. which i contributed to
 while i was employed
 (6) h. right
 (7) j. from the state of NJ mutual fund.
 and I'm entitled to a lump sum
 settlement which would be between
 16,800 and 17,800, or a lesser life
 annuity. and the choices of the annuity
 um would be \$125.45 per month.
 That would be the maximum
 with no beneficiaries
 (8) h. You can stop right there:
 take your money.
 (9) j. take the money.
 (10) h. absolutely.
 YOU'RE ONLY GETTING 1500 A YEAR.
 at 17,000, no trouble at all to
 get 10 percent on 17,000 bucks.

Harry interrupts her at (8) since he believes he has enough information to suggest a course of action, and tells her *take your money*. To provide **SUPPORT** for this course of action he produces an inference that follows from what she has told him in (7), namely *You're only getting 1500 (dollars) a year*. **SUPPORT** is a general relation that holds between beliefs and intentions in this model.

Presumably Jane would have no trouble calculating that \$125.45 a month for 12 months amounts to a little over \$1500 a year, and thus can easily accept this statement that is intended to provide the necessary **SUPPORT** relation, ie. the juxtaposition of this fact against the advice to *take the money* conveys that the fact that she is only getting 1500 dollars a year is a reason for her to adopt the goal of taking the money, although this is not explicitly stated.

5 Evidence of Acceptance

In section 3, I examine the assumption that: (2) Propositions representing beliefs and intentions get added to the context by the unilateral action of one conversant. I suggested that this assumption can be replaced by adopting a model in which agents' behavior provides evidence for whether or not mutual understanding has been achieved. I also discussed some of the effects of resource bounds, ie. cases of ensuring that or providing evidence that certain inferences dependent on what is

said are made.

Achieving understanding and compensating for resource bounds are issues for a model of dialogue whether or not agents are autonomous. But agents' autonomy means there are a number of other reasons why A's utterance to B conveying a proposition p might not achieve its intended effect: (1) p may not cohere with B's beliefs, (2) B may not think that p is relevant, (3) B may believe that p does not contribute to the common goal, (4) B may prefer doing or believing some q where p is mutually exclusive with q , (5) If p is about an action, B may want to partially modify p with additional constraints about how, or when p .

Therefore it is important to distinguish an agent actually **ACCEPTING** the belief that p or intending to perform an action described by p from merely understanding that p was conveyed. Other accounts legislate that helpful agents should adopt other's beliefs and intentions or that acceptance depends on whether or not the agent previously believed $\neg p$ [11, 8]. But agents can decide whether as well as how to revise their beliefs [4].

Evidence of acceptance may be given explicitly, but acceptance can be inferred in some dialogue situations via the operation of a simple principle of cooperative dialogue⁵:

COLLABORATIVE PRINCIPLE: Conversants must provide evidence of a detected discrepancy in belief as soon as possible.

This principle claims that evidence of conflict should be made apparent in order to keep **default** inferences about acceptance or understanding from going through. IRU's such as **PROMPTS**, **REPETITIONS**, **PARAPHRASES**, and making an **INFERENCE** explicit cannot function as evidence for conflicts in beliefs or intentions via their propositional content since they are informationally redundant. If they are realized with question intonation, the inference of acceptance is blocked.

In the dialogue below between Harry (h) and Ruth (r), Ruth in (39), first ensures that she understood Harry correctly, and then provides explicit evidence of non-acceptance in (41), based on her autonomous preferences about how her money is invested. .

Example 5:

- (38) h. and I'd like 15 thousand in a
2 and a half year certificate
(39) r. the full 15 in a 2 and a half?
(40) h. that's correct
(41) r. GEE. NOT AT MY AGE

In the following example, Joe in (14) makes a statement that provides propositional content that conflicts with

Harry's statement in (13) and thus provides evidence of non-acceptance.

Example 6

- (13) h. and -- there's no reason why you
shouldn't have an I R A for last year
(14) j. WELL I THOUGHT THEY JUST STARTED
THIS YEAR

Joe's statement is based on his prior beliefs. In both of these cases this evidence for conflict is given immediately. However when there is no evidence to the contrary⁶, and goals of the discourse require achievement of acceptance, inferences about acceptance are licensed as **default**. They can be defeated later by stronger evidence.

Without this principle, a conversant might not bring up an objection until much later in the conversation, at which point the relevant belief and some inferences following from that belief will have been added to the common ground as **defaults**. The result of this is that the retraction of that belief results in many beliefs being revised. The operation of this principle helps conversants avoid replanning resulting from inconsistency in beliefs, and thus provides a way to manage the augmentation of the common ground efficiently.

6 Other hypotheses

The first point to note is that the examples here are only a subset of the types of IRU's that occur in dialogues. I use the term antecedent to refer to the most recent utterance which should have added the proposition to the context. This paper has mainly focused on cases where the IRU: (1) is adjacent to its antecedent, rather than remote; (2) realizes a proposition whose antecedent was said by another conversant, (3) has only one antecedent. It is with respect to this subset of the data that the alternate hypotheses are examined.

A distributional analysis of a subset of the corpus (171 IRU's from 24 dialogues consisting of 976 turns), on the relation of an IRU to its antecedent and the context, shows that 35% of the tokens occur remotely from their antecedents, that 32% have more than one antecedent, that 48% consist of the speaker repeating something that he said before and 52% consist of the speaker repeating something that the other conversant said. So the data that this paper focuses on accounts for about 30% of the data.

⁶This displaying of evidence to the contrary was called an interruption in [15].

⁵This is a simplification of the COLLABORATIVE PLANNING PRINCIPLES described in [15].

6.1 Indirect Question Hypothesis

In example (1) of section 1, an alternative account of Ray’s repetition in (8) is that it is a question of some kind. This raises a number of issues: (1) Why doesn’t it have the form of a question?, (2) What is it a question about?, and (3) Why is it never denied?.

Of 171 IRU’s, only 28 are realized with rising question intonation. Of these 28, 6 are actually redundant questions with question syntax, and 14 are followed by affirmations.

If these are generally questions, then one possible answer to what the question is about is that Ray is questioning whether he actually heard properly. But then why doesn’t he use an intonational contour that conveys this fact as Ruth does in example 5? On an efficiency argument, it is hard to imagine that it would have cost Ray any more effort to have done so.

Finally, if it were a question it would seem that it should have more than one answer. While 50 of these IRU’s are followed by an affirmation such as *that’s correct, right, yup*, none of them are ever followed by a denial of their content. It seems an odd question that only has one answer.

6.2 Dead Air Hypothesis

Another hypothesis is that IRU’s result from the radio talk show environment in which silence is not tolerated. So agents produce IRU’s because they cannot think of anything else to say but feel as though they must say something.

The first point to note is that IRU’s actually occur in dialogues that aren’t on the radio[1]. The second question is why an agent would produce an IRU, rather than some other trivial statement such as *I didn’t know that*. Third, why don’t these utterance correlate with typical stalling behavior such as false starts, pauses, and filled pauses such as *uhhh*.

The dead air hypothesis would seem to rely on an assumption that at unpredictable intervals, agents just can’t think very well. My claim is that IRU’s are related to goals, that they support inferencing and address assumptions underlying mutual beliefs, ie. they are not random. In order to prove this it must be possible to test the hypothesis that it is only **important** propositions that get repeated, paraphrased or made explicit. This can be based on analyzing when the information that is repeated has been specifically requested, such as in the caller’s opening question or by a request for information from Harry. It should also be possible to test whether the IRU realizes a proposition that plays a role in the final plan that Harry and the caller negotiate. However this type of strong evidence against the dead air hypothesis is left to future work.

7 Discussion

It should be apparent from the account that the types of utterances examined here are not really redundant. The reason that many models of belief transfer in dialogue would characterize them as redundant follows from a combination of facts: (1) The representation of belief in these models has been binary; (2) The effects of utterance actions are either assumed to always hold, or to hold as defaults unless the listener already believed otherwise. This means that these accounts cannot represent the fact that a belief must be supported by some kind of evidence and that the evidence may be stronger or weaker. It also follows from (2) that these models assume that agents are not autonomous, or at least do not have control over their own mental states. But belief revision is surely an autonomous process; agents can choose whether to accept a new belief or revise old beliefs[4, 8].

The occurrence of IRU’s in dialogue has many ramifications for a model of dialogue. Accounting for IRU’s has two direct effects on a dialogue model. First it requires a model of mutual beliefs that specifies how mutual beliefs are inferred and how some mutual beliefs can be as weak as mutual suppositions. One function of IRU’s is to address the assumptions on which mutual beliefs are based. Second the assumption that propositions representing beliefs and intentions get added to the context by the unilateral action of one conversant must be dropped. This account replaces that assumption with a model in which the evidence of the hearer must be considered to establish mutual beliefs. The claim here is that both understanding and acceptance are monitored. The model outlined here can be used for different types of dialogue, including dialogues in which agents are constructing mutual beliefs to support future action by them jointly or alone.

How and when agents decide to augment the strength of evidence for a belief has not been addressed in this work as yet. Future work includes analyzing the corpus with respect to whether the IRU plays a role in the final plan that is negotiated between the conversants.

8 Acknowledgements

Discussions with Aravind Joshi, Ellen Prince and Bonnie Webber have been extremely helpful in the development of these ideas. In addition I would like to thank Herb Clark, Sharon Cote, Julia Galliers, Ellen Germain, Beth Ann Hockey, Megan Moser, Hideyuki Nakashima, Owen Rambow, Craig Roberts, Phil Stenton, and Steve Whittaker for the influence of their ideas and for useful discussions.

References

- [1] Jean C. Carletta. *Risk Taking and Recovery in Task-Oriented Dialogue*. PhD thesis, Edinburgh University, 1992.
- [2] Herbert H. Clark and Catherine R. Marshall. Definite reference and mutual knowledge. In Joshi, Webber, and Sag, editors, *Elements of Discourse Understanding*, pages 10–63. CUP, Cambridge, 1981.
- [3] Herbert H. Clark and Edward F. Schaefer. Contributing to discourse. *Cognitive Science*, 13:259–294, 1989.
- [4] Julia R. Galliers. Cooperative interaction as strategic belief revision. In M.S. Deen, editor, *Cooperating Knowledge Based Systems*, page 1. Springer Verlag, 1991.
- [5] H. P. Grice. *William James Lectures*. 1967.
- [6] Barbara J. Grosz, Aravind K. Joshi, and Scott Weinstein. Towards a computational theory of discourse interpretation. Unpublished Manuscript, 1986.
- [7] Barbara J. Grosz and Candace L. Sidner. Attentions, intentions and the structure of discourse. *Computational Linguistics*, 12:175–204, 1986.
- [8] Barbara J. Grosz and Candace L. Sidner. Plans for discourse. In *Cohen, Morgan and Pollack, eds. Intentions in Communication*, MIT Press, 1990.
- [9] Julia Hirschberg. *A Theory of Scalar Implicature*. PhD thesis, University of Pennsylvania, Computer and Information Science, 1985.
- [10] David Lewis. *Convention*. Harvard University Press, 1969.
- [11] Diane Litman and James Allen. Recognizing and relating discourse intentions and task-oriented plans. In *Cohen, Morgan and Pollack, eds. Intentions in Communication*, MIT Press, 1990.
- [12] Martha Pollack, Julia Hirschberg, and Bonnie Webber. User participation in the reasoning process of expert systems. In *AAAI82*, 1982.
- [13] Ellen F. Prince. On the function of existential presupposition in discourse. In *Papers from 14th Regional Meeting*. CLS, Chicago, IL, 1978.
- [14] Robert C. Stalnaker. Assertion. In Peter Cole, editor, *Syntax and Semantics, Volume 9: Pragmatics*, pages 315–332. Academic Press, 1978.
- [15] 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.