# Discourse Context

Richmond H. Thomason and Johanna D. Moore
Intelligent Systems Program
University of Pittsburgh
Pittsburgh, PA 15260
*Email:* thomason@isp.pitt.edu

Computer Science Department
University of Pittsburgh
Pittsburgh, PA 15260
*Email:* jmoore@cs.pitt.edu

**Abstract**

We discuss the problems of developing a theory of context that applies to the phenomena of naturally occurring discourse and that might be useful in assisting with the problems that arise in generating and interpreting discourse. The theories of context that have recently emerged in AI treat context as a determinant of literal meaning. This idea does in fact fit some natural language phenomena. But we suggest (1) that the account is inadequate for important effects that occur in discourse, and moreover (2) in discourse, the mutuality of context provides important constraints on the dynamics of context.

We explore these ideas in the paper, and also describe an ongoing project that seeks to develop an empirically adequate and computationally useful approach to the problem of contextual effects on inference in discourse.

Readers, please note that this is a provisional draft that describes ongoing work.

## Old and new theories of context

In both the philosophical and computational literature, much has been made of the importance of context in natural language interpretation. The theories of context that have been developed in this connection (for instance, in (3), (3), (3), and (3)) have stressed the part it plays in determining the meaning of natural language expressions.

The recent work on context that is being carried out in AI has much in common with the earlier work in philosophical logic, though it does not seem to draw directly on this work. The two approaches share the basic idea that context is an interpretation-shifting mechanism.

David Kaplan's theory provides a good representative version of the basic idea. The theory is confined to the indexicals 'I', 'here', and 'now'. In this case, a context can be modeled as a triple consisting of a person, a place and a time; these values then determine the interpretation of the corresponding indexicals. In at least one way, Kaplan's work is more sophisticated than the work on context that is emerging in AI—this is the distinction between content and character that emerges from his explanation of how a sentence like 'I am here now' can be valid, while 'I might not have been here now' can also be true. Kaplan's analysis shows that there are systematic variations between the contextual indices that influence interpretation of expressions, and the modal indices that keep track of differences in the facts that determine the truth values of interpreted expressions. One of the chief accomplishments of Kaplan's work is a clear demarcation of these differences.[1]

## Context and natural language

The recent work in AI on context starts with general problems having to do with representation and reasoning which have no intrinsic connection to natural language. Applications to context sensitivity in natural language then become a specialization of the more general theory. We believe that this strategy is more appropriate and more promising than a narrowly linguistic one; we should expect there to be far-reaching commonality between linguistic contextuality and other forms of contextuality, and to profit from exploiting this commonality.

At the same time, we have to remember the complexity that arises in the languages that have evolved for purposes of human communication. In the case of a designed language, there is some reason to think that

[1] On the other hand, Kaplan's limitation of the inquiry to these three indexicals no doubt results in some oversimplification. The more recent work serves to remind us that there are more complex examples. Consider, for instance, an example from (3, p. 63): 'Every man is wearing a tie', where 'every man' is contextually restricted to the contents of a certain room. It seems that 'Tomorrow, every man will be wearing a tie' is actually ambiguous: it can mean either that each man currently in the room will be wearing a tie tomorrow, or that each man in the room tomorrow will be wearing a tie. Such examples cannot be accommodated using Kaplan's two-way distinction between content and character. A theory of the sort advocated by Guha and McCarthy, which allows for explicit representation of the contexts, can express the ambiguity. But the project of explaining just the ambiguities that natural language allows seems to be a challenging one.

nothing will be found that was not put there by the designer. But in the case of natural language, a theory has to go hand in hand with an empirical method of testing linguistic hypotheses, and these methods often reveal phenomena of astonishing depth and complexity. We will argue here that the evidence suggests that the theory of context as a determinant of literal interpretation has to be modified and extended to account for the phenomena that are exhibited by discourse in natural language.

Our discussion will concentrate on two factors that play a crucial role in discourse: effects of context on inference, or *nonliteral* interpretation, and the mutuality of context.

## Nonliteral meaning

There are difficulties with the view that contextual effects in natural language are confined to "K-effects," Kaplan-like effects on the literal meaning of a text. The purpose of this section is to explain this point to readers who are familiar only with the recent theories of context in AI.

To begin with, consider a clear case of a K-effect.

1a. *Demonstration:* [The speaker points at a book.]

1b. *Utterance:* A: Bring me that book.

Here, a preparatory action (1a) serves to alter the context. The subsequent utterance (1b) uses the altered context to issue an imperative.

Contrast this with a contextual "priming" effect on preferences for certain readings of an ambiguous expression. The sentence

2. *Text:* The astronomer married a star.

is difficult to interpret because its subject somehow creates a preference for the wrong reading of its object.

This priming phenomenon may seem superficially similar to example (1); you could say that the subject phrase of (2) chooses a context in which 'star' has a certain denotation, just as (1a) chooses a context in which 'that' refers to a certain book.

But virtually all theories of natural language semantics treat lexical ambiguity, such as the ambiguity between the two readings of 'star', in an entirely different way. The grammar of a semantically interpreted language will associate a set of readings or alternative interpretations with each expression; but the members of this set are not related in any systematic way to contexts. Notice that in example (2), it is possible to access the other reading of 'star' after recovering from the initial feeling of surprise. It is, of course, possible to talk about patterns of preferences or likelihoods for readings of ambiguous lexical items. But it isn't clear what such patterns have to do with the theories of context inspired by logic. These theories all seem to assume that lexical ambiguities have been resolved in advance.

The inferential process of selecting a preferred reading of an ambiguous expression has not in fact been much studied by the semanticists. Although a narrowly linguistic theory might do well to ignore this phenomenon, it is vitally important in modeling the process of actual discourse interpretation; ambiguity is perhaps the single greatest technical problem in this area of computational linguistics.

Equally important in actual discourse are processes that infer additional, nonliteral meaning. Gricean implicatures[2] are such inferences, and in fact one of Grice's tests for distinguishing implicatures from literal meanings is sensitivity to context.[3] But Grice is very careful to point out that implicatures are not part of the literal meaning of the utterance. If Grice is right about this, there are vitally important effects of context in discourse that do not have to do with variation of literal meaning over the space of contexts.

Consider this point in relation to an example—say, the scalar implicature in 'The oven is warm'.[4]

3a. *Utterance:* A: Is the oven hot?

3b. *Utterance:* B: It's warm.

3c. *Inference:* The oven isn't hot.

4a. *Utterance:* A: Is the oven warm?

4b. *Utterance:* B: It's warm.

4c. *Inference:* No inference about whether the oven is hot is warranted.

In the context of (3a), (3b) invites the inference that the oven isn't hot. Because this inference is cancelled in the context of (4a), Grice's contextual cancelability test indicates that this is an implicature.

Obviously, though, this test will not distinguish implicatures from K-effects; contextuality is the hallmark of a K-effect. But consider the following example, which invokes Grice's *explicit cancelability* test. Here, the inference is blocked not implicitly by the context, but by an explicitly expressed denial.

5a. *Utterance:* A: Is the oven hot?

5b. *Utterance:* B: It's warm. I'm not sure if it's hot.

Explicit cancelability does in fact seem to distinguish implicatures from cases where there is a good case to be made for K-effects of context.[5]

[2]See (3).

[3]See (3, p. 44).

[4]For information on scalar implicature, see (3) and (3).

[5]Note the contrast with K-effects. It is quite impossible to explicitly cancel a demonstration; for instance, you can't

The purpose of Grice's tests is to set limits to the potential effects of context on literal meaning; Grice feels that it is implausible to assume here there is any context in which 'It's warm' literally means that the oven is warm but not hot. Underlying Grice's desire to limit the effects of context on literal meaning is a prohibition against multiplying the senses of a word[6] beyond necessity. In a more contemporary vein, we could argue that treating these inferences as matters of literal meaning would hopelessly complicate the lexicon.

The purpose of Grice's theory of implicature is to explain why it is reasonable to make certain inferences even when they are not part of the literal meaning. Unfortunately, this theory is much too vague and impressionistic to provide explanations of examples at a satisfactory level of detail. However plausible Grice's general argument may be (and it is very plausible) it remains inconclusive in the absence of a really successful positive account of the inferences that support implicature.

In a series of recent papers,[7] Nicholas Asher, Alex Lascarides, and Jon Oberlander have argued for a treatment of such inferences that is based on commonsense entailment, a system of nonmonotonic logic. According to this approach, these inferences are the nonmonotonic consequences of certain discourse postulates and the literal interpretations of the sentences in a text. The inferences that can be obtained from a sentence embedded in a larger text are contextually conditioned by the interpretations of the surrounding sentences. These studies concentrate on the problem of *inferred discourse relations*; e.g., whether in a case like the following example, the event indicated by (6b) follows (6a), or is a cause of the event indicated by (6a) and precedes it.

6a. *Text:* Max fell.

6b. *Text:* John pushed him.

The general idea of a textually conditioned context which affects the appropriate inferences is close to the one that we will explore in this paper. We advocate a different approach to the reasoning, in part because we believe in using techniques that have already proved themselves in natural language processing, and in part because we feel that the account we will offer may be better suited to the phenomena of task-oriented dialogue.

The experience of developing the commonsense entailment account of discourse relations indicates that if nonmonotonic logics are used for this task they will probably need a more flexible and powerful way of managing priorities than is presently available. At this early stage in the formalization of discourse inferences, given the evident complexity of the problems, we believe that it is good to simultaneously develop many approaches. The abductive framework we will explore later in this paper is in fact a close relation to the family of nonmonotonic logics.[8]

To summarize where we have gotten to at this point:

1. Along with most linguists and many philosophers of language, we advocate a Gricean distinction between the semantic and the pragmatic effects of context. The former influence the literal interpretation of expressions, while the latter influence the inferences that are appropriate.
2. Unfortunately, we lack a satisfactory general theory of the pragmatic effects of context.
3. However, it may be possible to use a nonmonotonic framework to formalize the pragmatic effects.

To these points, we now add the thought that the needs of communication imply that the context used in discourse must be mutually accessible to speaker and hearer.

## The mutuality of context

A lawyer might infer from a witness's statement that the witness is concealing relevant evidence. Inferences of this sort, which—however well supported they may be—are not intended by the speaker and not recognized as such by the hearer, are not the target of our study. We are interested only in inferences that are meant.

Assuming that these inferences depend on the context of the utterance, the context must be mutual for conversation to remain coherent. We assume (along with theoreticians of discourse like Herbert Clark[9]) that what is meant must be mutually known.[10] So, if what is meant depends on what is intended to be inferred, and what is intended to be inferred depends on the context, the speaker's and the hearer's views of the context must coincide (and must be supposed to coincide).

This constrains the sort of context that is needed for discourse inference: it should only contain things that can be mutually known to the participants. Although this constraint doesn't tell us what a discourse context is, or how it serves to guide discourse reasoning, it tells us a good deal about how a discourse context is maintained.[11]

point at a book and say 'Bring me that book; but I don't mean the book I'm pointing at, I mean the other book'.

[6] We take this to extend to the context-dependent meanings of a word as well.

[7] See (3), (3), (3), and (3).

[8] See (3).

[9] See (3).

[10] Or, more accurately, mutually supposed for the sake of the conversation. See (3).

[11] The following account is based on ideas from (3) and

A mutual initial discourse context can only be created if the participants have a good sense of what they are likely to mutually assume at the beginning of a conversation. Studies such as those reported in (3, Chapters 1–3) show that people indeed have a robust sense of the common ground that is appropriate for generating and interpreting referring expressions. This common ground not only draws on general background that any participants will normally share, but may exploit more particular material that they share as colleagues or close acquaintances. Also it may use immediate perceptual information that is mutually available to the participants; this last feature is one source of the difference between telephone and face-to-face conversations. Conversational participants tailor the common ground to their audience; the exploitation of the richer common ground that is shared by close acquaintances can make it much more difficult for an overhearer to follow a conversation than for the person for whom the utterance is intended.[12]

The discourse context evolves in the course of a conversation. Stalnaker's work concentrates on the simplest update mechanism. He argues from phenomena such as presupposition that the discourse context contains a record of what has been established up to a given stage of the conversation, and that a successful assertion will then add the asserted proposition to this record. The mechanism is generalized in (3) to include many other parameters. These serve to keep track of definite reference, verbs like 'come' that incorporate speaker location, relative modality (e.g., the relative strictness with which a word like 'can' is interpreted), information from evolving plans, and other sources of information.

It follows from the mutuality constraint that contextual dynamics must be mutual; otherwise, even though the participants begin a conversation with the same discourse context in view, this initial coordination will be lost as the conversation evolves. For the case of assertion, this mutual dynamics is modeled in (3) by using mutual defaults (default rules that are mutually known to the conversational participants) to drive the evolution of the discourse context.

This idea can be extended from assertion to speech acts in general. The direct purpose of a speech act is to affect some parameter in the discourse context. To account for proposals, for example, we would include a component in the discourse context to record the proposed course of action that is "on the table." A proposal to follow $\alpha$, where $\alpha$ is a course of action, will have the effect of placing $\alpha$ in the proposal component of the discourse context.[13] Obviously, the proposal component can be modeled as a fluent, while the discourse context itself is treated as a situation or a distinguished part of the situation (a part that is mutually known). This approach has the advantage of providing a direct link to planning formalisms. There also is a connection to the SharedPlans formalism of (3).

The mutuality constraint has at least one major consequence in applying these ideas to the design of systems for natural language processing. The exploitation of a mutual discourse context depends on the attunement of the processes of generation and interpretation. Assuming, with Stalnaker and Lewis, that assertions are added to the conversational record, for coordination to be preserved it will be necessary for the speaker's and the hearer's view of what is conveyed by an utterance to coincide. Thus, the generator needs to take into account the interpretive process, to ensure that the utterances it produces will have the desired effect; similarly, the interpreter needs a model of the intentions that can be supposed to underlie an utterance. To do justice to discourse context, work in natural language processing needs to take into explicit account both generation and interpretation, so that the interactions between the two processes can be investigated. We stress this because relatively little work has tried to do this yet.

Though we believe that we will ultimately need a common framework for generation and interpretation in order to obtain an adequate account of the phenomena, we will not attempt to develop that framework in this paper. Instead, we will show how a planning approach to generation, such as that of (3), can be used to produce utterances that take into account the hearer's ability to fill in detail through discourse inference.[14] In the last section, we will briefly describe the abductive framework in which we hope to provide a more integrated and general account.

## Testing and applying models of discourse context

In applying these ideas to natural language processing, the major difficulty that needs to be overcome is the enormous complexity of the factors that can be brought to bear in unconstrained, general purpose conversation. We are mitigating this problem by working in a task-oriented discourse domain in which the linguistic and nonlinguistic reasoning are limited. In this

subsequent writings of Robert Stalnaker's, from (3), from (3), and from (3).

[12] See (3, chapter 8).

[13] Obviously, there has to be some mechanism for using the discourse context for general purpose reasoning. In particular, assertions may be adopted as individual beliefs, and proposals may become individual intentions. But these mechanisms may depend on factors such as trust, and are less a matter of discourse than of general inference from perceptual inputs.

[14] Without a common framework for generation and interpretation, we have to give the generator indirect knowledge of the interpretive process, through a set of "discourse rules." These rules could be justified if an explicit model of the interpretive process were provided.

limited domain, we can test how human subjects maintain coordination in their discourse, and can seek to model their performance by simulating generation and interpretation mechanisms that make strong use of the mutual context. This provides an empirical approach to the discourse phenomena,[15] as well as some hope that a simulation can be implemented.

For this domain, we have adapted the design reported by Walker in (3). Using a simulation approach to discourse to explore hypotheses about redundancy and constraints on the memory capacity of discussants, she designed a "collaborative exploratory planning" task in which two discussants are provided with some capital and inventories of available furniture. Neither agent knows what furniture is available to the other. Together, the agents are supposed to buy furniture for a hypothetical apartment. Walker uses a simple sum of the values of the furniture items to represent the success of the discourse.

We have complicated Walker's task somewhat to allow for slightly more complex conversations.[16] The following model of the discourse phenomena is based on preliminary investigations with human subjects that we have been carrying out over the last six months.

We are not yet in a position to present a formalization of the account of context and discourse planning that we are building to model generation in this domain. Here, we only point out that the following ideas can be formalized using the Situation Calculus. In this formalization, situations correspond to states of the discourse context. Typical operators correspond to speech acts, such as proposals, objections, and assertions; other operators correspond to more strategic actions, such as supporting a proposal.

The domain reasoning in this task consists of exploratory search of the plan space that is induced by a single PUT-FURNITURE operator, which places a unit of furniture in a room. Discourse reasoning is induced by domain reasoning in two ways: (1) by the need for mutual agreement on the domain plans (this need is represented as a precondition on the PUT-FURNITURE operator), and (2) by the need to satisfy knowledge preconditions of the PUT-FURNITURE operator. Discourse reasoning can also be induced by discourse; for instance, a question can induce a goal to provide an answer.

Humans performing this task will explore the domain plan space by incrementally agreeing to add items of furniture until their money is spent. One piece of the discourse context, then, is the mutually agreed-on current partial domain plan: the Mutual-Plan component. This part of the context would affect the reference of indexicals like 'next'; as we will see, it also plays an important part in guiding discourse inference.

The previous discourse also has to be included in the discourse context in order to account for anaphoric pronouns and similar phenomena. To model proposals in this domain, we incorporate an "On-Table-Proposal" component in the context. This component is occupied by whatever domain action has been proposed in the previous turn; if no action has been proposed, it is empty. The contents of the On-Table-Proposal component are added to the Mutual-Plan component on the next discourse turn unless an objection to the proposal is registered on this turn; in that case, the component will either be emptied or occupied with a counterproposal. In this domain, the conversants share high-level goals, a mutual planning strategy, a position in the domain plan space, and some knowledge about the available furniture. This shared information yields a mutual notion of the reasonable options that are open on a given turn; these salient options we also suppose to be part of the context.

We illustrate the role that context plays in generation by considering how in a certain context in our discourse task an appropriately brief utterance might be generated, rather than a longer and relatively decontextualized version.

First, some stage setting.

We assume that the players have adopted a domain planning strategy of working on one room at a time.[17] They share knowledge about what furniture is appropriate for the various rooms. The players A and B are currently working on the living room.

At the previous stage of the game B has made the following utterance.

7. *Utterance:* B: I have a red couch for $400. Let's put it in the living room.

We assume that the utterance has been appropriately interpreted and the discourse context has been updated accordingly. In particular, this adds

put(couch1, LR)

to the On-Table-Proposal component of the discourse context.

It is A's turn to respond. A has a cheaper couch, one that B doesn't know about, and wishes to reject B's proposal, substituting the new couch. We want to explain how in this context, A can generate the following utterance in the expectation that it will be understood.

8. *Utterance:* A: I've got a blue one for $300.

Utterance (8) is an abbreviated version of the following expanded sequence of utterances.

[15] See (3).

[16] Essentially, we have used a more complex scoring system that is supposed to measure the utility of the furniture that is purchased, and we have departed from using a rigid script to guide the discourse.

[17] Some, but not all of our subjects seem to follow this strategy.

9a. *Utterance:* A: No, let's not put the red couch for $400 in the living room.

9b. *Utterance:* A: I have a couch.

9c. *Utterance:* A: It costs $300.

9d. *Utterance:* A: $300 is less than $400.

9e. *Utterance:* A: So let's put my couch in the living room.

Here is a brief explanation of how the context plays a part in justifying this abbreviation.

Saying just (8) implies the rejection (9.a), because extending the current partial domain plan by putting another couch in the living room is a competing option in this context—that is, it is a different reasonable option. (We assume as a rule of discourse that proposing an option that competes with the current proposal implies a rejection of that proposal. This rule, like all rules of discourse, has to be mutually known.)

Reporting that you have a piece of furniture makes the new piece part of the common ground. This addition may then alter the context by adding a new option to the set of current mutual options. We assume as a rule of discourse that if the new option that is introduced in this way is a reasonable option in the context, then the utterance that introduces it has the effect of a proposal. So the proposal (9.e) can be inferred from (8).

The eliminability of (9.d) follows from considerations that are relatively context-free.

On the other hand, (9.b) and (9.c) are needed to provide support and knowledge preconditions for the high-level goal of making a counterproposal.

Once this reasoning has determined what can be eliminated from the utterance, considerations of the sort treated in (3) and (3) suffice to generate the wording of (8).

## An abductive model of reasoning from context

The application of weighted abduction to natural language interpretation is described in (3); in its simplest form, the idea is to identify the abductive explanation of the logical form of an utterance with the inferences that constitute its interpretation. A more general version of the approach integrates this interpretive process with parsing of the surface utterance.[18] In subsequent work, which is mainly unpublished, Jerry Hobbs has discussed how the approach can be used to model generation as well as abduction.[19]

The reasoning method itself is described in (3, Section 4); also consult (3). This is a method of resolution theorem proving for Horn clauses in which assumptions can be used in proofs, as well as axioms. Assumability is assigned a prior cost; the prior cost of an assumption is adjusted by the abduction algorithm in a context where previous assumptions have been made to the extent that the assumption unifies with axioms and hypotheses already assumed.

In this way the abductive approach provides an automatic way to convert linguistic contexts—where these contexts are realized as a series of assumptions—to inferences. Along the way, the method provides the nonmonotonic effects that are stressed by Asher, Lascarides, and Oberlander. Sentences that have already been processed provide abductive hypotheses. These hypotheses then serve as axioms for further abductions, and so reduce the cost of anything that unifies with them to zero. In this way, the context will affect the hypotheses that are preferred by the abduction algorithm.

For example, this approach accounts for textual effects on interpretation by requiring the interpretation process to explain not only the interpretations of single sentences, but the coherence of adjacent subtexts. The following inference, for instance, could be obtained by postulating that the enablement relation renders adjacent segments coherent; of course, common-sense axioms providing enablement relations between specific act types are also needed.[20]

10a. *Text:* Laura picked up a screwdriver. She unscrewed the cover.

10b. *Inference:* Laura used the screwdriver to unscrew the cover.

So far as Hobbs has developed it, this theory does not yet involve an independent theory of context, and only takes the effects of context into account implicitly, through the effects of the hypotheses abduced from surrounding text on the abduction process. (Recall that the conclusions that are reached by the weighted abduction algorithm are affected by hypotheses that have already been obtained, since these affect the costs of further hypotheses.) The task-oriented dialogue application that we discussed in the previous section, however, requires an explicit model of context, similar in many ways to those that have been emerging in the recent AI literature. This can easily be incorporated

[18] See (3, pp. 104–108).

[19] In general, planning can be carried out in an abductive framework by treating it as the problem of explaining the goal, in a reasoning context in which relatively low costs are assigned to actions that are under the control of the agent. In its simplest form, generation is the problem of explaining

$$\exists s\, Expresses(s, p),$$

where *Expresses* is the semantic relation between a sentence and the proposition it expresses. We are currently investigating with Hobbs ways in which to bring an explicit representation of the discourse context into this model, and to take the interpretive process into account.

[20] See (3, Section 6.3) for a detailed account of similar examples.

into the abductive framework by making the relation

$$Expresses(s, p)$$

sensitive to context, replacing it by

$$Expresses(s, p, c).$$

As we said, we are currently exploring this idea in collaboration with Hobbs.

## Acknowledgments

This material is based on work supported by the National Science Foundation under Grant No. IRI-9314961, "Integrated techniques for natural language generation and interpretation."

## References

[1] Douglas Appelt. *Planning English Sentences.* Cambridge University Press, Cambridge, England, 1985.

[2] Gerhard Brewka and Kurt Konolige. An abductive framework for general logic programs and other non-monotonic systems. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, pages 9–15, Morgan Kaufmann, Los Altos, California, 1993.

[3] Herbert Clark. *Arenas of Language Use.* University of Chicago Press, Chicago, 1992.

[4] Robert Dale. *Generating Referring Expressions: Constructing Descriptions in a Domain of Objects and Processes.* MIT Press, 1992.

[5] Gerald Gazdar. *Pragmatics: Implicature, Presupposition, and Logical Form.* Academic Press, New York, 1979.

[6] H. Paul Grice. *Studies in the Way of Words.* Harvard University Press, Cambridge, Massachusetts, 1989.

[7] Barbara Grosz and Sarit Kraus. Collaborative plans for group activities. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, pages 55–63, Morgan Kaufmann, Los Altos, California, 1993.

[8] Ramanathan Guha. *Contexts: a Formalization and Some Applications.* Technical Report STAN-CS-91-1399, Stanford Computer Science Department, Stanford, California, 1991.

[9] Julia Hirschberg. *A Theory of Scalar Implicature.* Garland, New York, 1991.

[10] Jerry Hobbs, Mark Stickel, Douglas Appelt, and Paul Martin. Interpretation as abduction. *Artificial Intelligence*, 63:69–142, 1993.

[11] Pamela Jordan and Richmond Thomason. Empirical methods in discourse: limits and prospects. In *1995 Workshop on Emipirical Methods in Discourse*, AAAI, 1995.

[12] David Kaplan. On the logic of demonstratives. *Journal of Philosophical Logic*, 8:81–98, 1978.

[13] Alex Lascarides and Nicholas Asher. Discourse relations and common sense entailment. In *Proceedings of the 29th Meeting of the Association for Computational Linguistics*, pages 55–63, Association for Computational Linguistics, 1991.

[14] Alex Lascarides and Nicholas Asher. Lexical disambiguation in a discourse context. 1993. Unpublished Manuscript.

[15] Alex Lascarides and Nicholas Asher. Temporal interpretation, discourse relations and common sense entailment. *Linguistics and Philosophy*, 16:437–494, 1993.

[16] Alex Lascarides and Jon Oberlander. Temporal coherence and defeasible knowledge. In *Proceedings*

*of the Workshop on Discourse Coherence*, Edinburgh, 1991.

[17] David Lewis. Scorekeeping in a language game. *Journal of Philosophical Logic*, 9:339–359, 1979.

[18] John McCarthy. Notes on formalizing context. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, pages 81–98, 1993.

[19] Johanna Moore. *Participating in Explanatory Dialogues.* MIT University Press, 1995.

[20] Robert Stalnaker. Presuppositions. *Journal of Philosophical Logic*, 447–457, 1973.

[21] Mark Stickel. Rationale and methods for abductive reasoning in natural-language interpretation. In R. Struder, editor, *Notes in Artificial Intelligence 48*, Springer, Berlin, 1989.

[22] Richmond Thomason. Accommodation, meaning, and implicature: interdisciplinary foundations for pragmatics. In P. Cohen, J. Morgan, and M. Pollack, editors, *Intentions in Communication*, MIT Press, Cambridge, Massachusetts, 1990.

[23] Richmond Thomason. Propagating epistemic coordination through mutual defaults i. In R. Parikh, editor, *Proceedings of the Third Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 29–39, Morgan Kaufmann, Los Altos, California, 1990.

[24] Marilyn Walker. *Informational Redundancy and Resource Bounds in Dialog.* PhD thesis, University of Pennsylvania, 1993.