

Enlightened Update: Informative Presupposition in a Computational Architecture for Pragmatic Reasoning*

Richmond H. Thomason
University of Michigan

Matthew Stone
Rutgers University

David DeVault
University of Southern California

Abstract We relate the theory of presupposition accommodation to a computational framework for reasoning in conversation. We characterize the evolution of the conversational state as following from two key principles: that interpretation is a form of intention recognition; and that intentions are complex informational structures, which specify commitments to conditions and to outcomes as well as to actions. Presuppositions, on this view, are private commitments the speaker makes in using an utterance but expects the listener to recognize based on mutual information. A natural consequence of these principles, we argue, is that conversation can move forward not just through the positive effects of interlocutors' utterances but also from the retrospective insight interlocutors gain about one another's mental states from observing what they say and do. In particular, we identify a specific phenomenon we call ENLIGHTENED UPDATE, in which the conversational state is updated to reflect the speaker's recognizable private commitments as well as the explicit contribution of the utterance. Enlightened update can be seen as a special case of presupposition accommodation, in that the hearer accommodates, or comes to accept as new information, what the speaker

* This paper is a revision of a manuscript originally prepared for a workshop on presupposition accommodation at the Ohio State University in Fall 2006. We gratefully acknowledge the organizers of the meeting, Craige Roberts, Donna Byron and Scott Schwenter, and the grant that supported it, NSF BCS-0548305, for the chance to develop and present this work, and Matthew Stone's Visiting Fellowship at the University of Edinburgh in Academic Year 2005-2006, supported by the Leverhulme Trust, during which the original manuscript was prepared. The present version benefits from the feedback from the workshop organizers, from our commentators at the workshop, Mats Rooth and Steffan Larsson, and from other participants there, and from comments by Luciana Benotti, Alex Lascarides, Iris Oved, Chung-chieh Shan, and David Traum. This work was sponsored in part by the U.S. Army Research, Development, and Engineering Command (RDECOM), and in part by NSF HLC-0308121, CCF-0541185, HSD-0624191, and IIS-1017811. Statements and opinions expressed do not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

presupposes, or commits to in using the utterance, but it is a side effect of interlocutors' ordinary reasoning in conversation, not a special strategy for adjusting, repairing or co-operating on the state of the conversation. We formalize and implement these ideas for a simple conversational agent, and draw on this case study to argue that pragmatic reasoning is holistic in character and continuous with common-sense reasoning about collaborative activities, and show that it can be characterized by associating specific, reliable interpretive constraints directly with grammatical forms.

Keywords: presupposition, accommodation, coordination, collaboration, intention

1 Introduction

Theorizing in pragmatics is complicated by the rationality that interlocutors bring to bear in their interactions with one another. In formulating utterances, speakers can use the obvious strategic choices they make to get across content beyond what those utterances conventionally mean (Grice 1989). In many cases, speakers can even count on their audience to recognize a hidden significance in utterances that might otherwise seem to violate norms (Grice 1989, Lewis 1979, Thomason 1990). In describing such interactions, researchers have inevitably postulated conflicting ways interlocutors can resolve their strategic dilemmas, through generalizations that come with caveats and exceptions. The theory of presupposition and accommodation is a case in point. Speakers' presuppositions are normally background information that they and their audience take for granted. Nevertheless, in many contexts, speakers succeed in getting new information across to their audience by using an utterance in a way that seems to presuppose it (Stalnaker 1975). In such contexts, it is now common to say that the presupposition is accommodated.

To develop a more systematic perspective on pragmatic phenomena, theorists must develop their ideas in a framework that relates pragmatic reasoning to reasoning in other kinds of social interaction, characterizes pragmatic reasoning explicitly, and allows predictions to be stated and tested precisely. This paper takes up this challenge. We formalize aspects of collaborative reasoning in conversation in a computational architecture, and use this architecture to characterize some of the reasoning that might underpin presupposition and accommodation.

Our work leads to a distinctive new way of understanding what presupposition is and how presupposition is recognized in conversation. This perspective explains many natural cases of informative presupposition as side effects of interlocutors' ordinary interpretive reasoning. In particular, we argue, these cases find nothing exceptional in the speaker's presupposing what they do, or in the audience's reasoning in understanding the speaker. Our new account acknowledges, however, that some cases of informative presuppositions do require exceptional interpretive effort. Our framework helps to delimit these cases, and to suggest what it is about these agents' circumstances, aims and choices that makes these cases special.

Our ultimate goal is to develop an explicit, computational model of the reasoning that allows interlocutors to work together in specific contexts to reach a shared understanding of their contributions to conversation. Such a model will specify the algorithms that accomplish this reasoning and the symbolic data structures over which the algorithms are defined. But this model will inevitably go beyond the computations themselves. It will, in addition, offer insight into the underlying PROBLEMS that interpretation algorithms need to solve, and into the real-world INFORMATION that pragmatic structures need to represent. As has been long recognized in Cognitive Science (Newell 1981, Marr 1982), it is at this "knowledge level", or level of "computational theory", that computational approaches can most directly engage with other methodologies for the study of mind and behavior. Accordingly, we begin this paper with a knowledge-level outline of the approach we will develop. We characterize conversation as a form of collaboration and characterize pragmatic reasoning in particular as an outgrowth of speakers' and listeners' collaborative efforts to coordinate on communicative intentions.

1.1 The requirements of conversational coordination

Like many others across Cognitive Science (see, for instance, Lewis 1969, Clark 1992), we view communication as a problem of coordination. In conversation, speakers produce each utterance with a particular interpretation in mind; listeners infer an interpretation for each utterance. In paradigm cases of successful communication, the interlocutors' interpretations of each utterance perfectly agree.

At the knowledge level, the problem of coordination places strong constraints on the information interlocutors should use in conversation; they should only use information that is public, or MUTUALLY SUPPOSED (Stalnaker

1975, Thomason 1990) in a sense we will describe more fully below. We summarize this by the coordination principle in (1).

- (1) In conversations, speakers must formulate, and listeners must infer, an interpretation for each utterance that is uniquely recognizable from the form of the utterance, using mutually supposed information.

No matter what content a theory of communication assigns to utterance interpretations, the theory must respect Principle (1). However, we take a specific choice about interpretation here. Again, like many others across Cognitive Science inspired by Grice (1957), we appeal to the theory of intention to flesh out the content of utterance interpretations.

1.2 Intentions

Interactions between Philosophy and Artificial Intelligence (AI) have been particularly fruitful in the area of rational agency. These ideas can be applied to conversation, viewed as the rational exchange of information. As in the theories of agency developed by Bratman (1987) and Pollack (1990, 1992), we understand an intention as a complex mental attitude that summarizes an agent's reason to act.

Intentions perform a range of functions in rational agency. Intentions help agents relate their actions to their goals in a way that facilitates problem solving, such as planning complex action sequences, coping with obstacles and opportunities that arise in a dynamic world, and understanding why actions fail and how to respond when they do. Of course, intentions can also help agents to work more effectively together. It is because intentions play these varied roles in agents' deliberations that researchers have analyzed intentions as complex informational structures. Intentions delineate the real-world circumstances that the agent is committed to, characterize the actions that the agent undertakes, and explain, in the context of the agent's theory of the world, why taking these actions in those circumstances will lead to an outcome that the agent wants.

From a Gricean perspective, the interpretation process in conversation is a species of INTENTION RECOGNITION—a hearer's reasoning task is to reconstruct the speaker's communicative intention. This, together with the above characterization of intention, yields the following characterization of

the interpretations that interlocutors infer in pragmatic reasoning.

- (2) Interpretations outline not just the content of a specific contribution to conversation, but also the speaker's underlying intention, including the speaker's commitments to the structure of the utterance (as analyzed by the grammar) and to the relevant extralinguistic facts that link grammatical meaning to conversational content.

Of particular note for our purposes is the fact that utterances, like other actions, may depend for their effects on preconditions about the situation in which they are performed. Thus, as made explicit in (2), a speaker's communicative intention will register apposite commitments about the conversational situation.

Models of intention thus suggest a knowledge-level account of conversation framed in terms of the common-sense information about utterances that is required in a communicative intention. Our formalism provides data structures—complex, structured pragmatic representations—that package this information; our model of pragmatic reasoning uses these representations to carry out the inferences involved in attributing a communicative intention to a speaker and calculating the import of recognized intentions for interlocutors' ongoing collaboration (Stone 2004a,b). Thus, in contrast to many approaches to formal pragmatics, we do not view our formalism as describing the INPUT to higher-level processes of reasoning about intentions and collaboration. It spells out a SUBSTRATE for this pragmatic reasoning and must itself be narrated in high-level common-sense terms.

1.3 Conversationally manifested private commitments

Our view, ultimately, is just that reasoning processes in conversation are continuous with reasoning processes in other kinds of coordinated, collaborative activity. This idea is familiar (Clark 1992)—but characterizing these processes in computational terms offers new kinds of precision, clarity, and insight into pragmatic phenomena that seem to depend on interlocutors' inferential abilities. We focus in this paper on the relationship among grammatical knowledge, world knowledge and mutual attitudes in presupposition.

We seek to clarify and understand this relationship through a detailed consideration of the interpretations that collaborative interlocutors can assign to their utterances. According to Principle (2), interpretations are complex structures, and can express commitment to propositions in many ways. As

an initial illustration of how a speaker's commitment to a proposition can be expressed, consider Stalnaker's (1975: 202) example (3):

(3) I have to pick up my cat at the veterinarian.

The best interpretation of (3) may be an intention to communicate the proposition that I have to pick up my cat at the veterinarian, an intention that contains a commitment to the proposition that I have a cat. In terms of mutual attitudes, this second commitment may well be private at the time of the utterance, but easily recognized from the utterance itself. Here the speaker expresses an intention to communicate one (private) proposition, and at the same time reveals, as part of this intention, a background commitment to another (private) proposition.

Our framework provides a general approach to the coordination needed in order for speakers to successfully communicate using their private background commitments. In particular, Principles (1) and (2) give a special place in pragmatic theory to PUBLICLY RECOGNIZABLE PRIVATE COMMITMENTS. Such commitments play the role in our theory that SPEAKER PRESUPPOSITIONS play in Stalnaker's approach to pragmatic presupposition (Stalnaker 1973, 1975), a point we return to in Section 2.3. In using an utterance, as articulated in Principle (2), a speaker commits FOR HERSELF that the conversational situation is as the grammar requires for her utterance to contribute the content she wants. As we shall see, the conversational situation may not encompass everything the speaker believes, but neither must it be limited to the interlocutors' public information. What matters in the first instance is just that the speaker herself believes that the conversational situation is how it needs to be, and commits to this belief. However, Principle (1) also applies. The speaker's commitment is itself PART of the interpretation, and the listener must recognize the interpretation. So according to Principle (1), the speaker must also make sure that this commitment can be recognized from the utterance and other public information.

We can highlight the significance of this characterization by observing that most utterances will require the speaker to commit in this way to MANY specific facts. Each fact might of course be public. But even commitment to information that is NOT public can easily be made recognizable. The speaker need only commit ELSEWHERE to enough public information to clue the listener in to the overall picture she has in mind. Principle (1) says only that the speaker must allow the hearer to infer one interpretation as the best, on the basis of mutual background information.

Our framework casts this coordination problem, of how interlocutors can expect each other to agree on the best interpretations for their utterances, as an abductive reasoning problem. We discuss this abductive reasoning problem in the next section.

1.4 Abductive reasoning and enlightened update in conversation

The knowledge level account we have presented follows naturally from the central idea that interpretation is a form of intention recognition, and that intentions are complex informational structures. But since it is meant to be a computationally useful account, it requires and must inspire an explicit, implementable model of the associated reasoning. Here, we rely especially on models of conversational reasoning developed by Jerry Hobbs and colleagues at SRI—models of interpretation as abduction (Hobbs, Stickel, Appelt & Martin 1993). Hobbs treats interpretation as a form of inference to the best explanation. His colleague Mark Stickel showed how this reasoning could be implemented in a logic programming framework, and Hobbs showed how the implementation could be used to model a broad sample of pragmatic phenomena.

According to the traditional characterization, abductive reasoning leads from evidence to a hypothesis that explains the observation. An abductive reasoner with a stock of beliefs will observe a proposition E , and add this proposition to its stock of beliefs along with a proposition H , where H is the best explanation of E . We can model interpretation as abduction by taking E to be the observation that an utterance has been made and taking H to specify the speaker's communicative intention in using that utterance. In particular, in keeping with Principle (2), we understand H to spell out the commitments the speaker has made about the initial conversational situation, the grammatical analysis of the utterance, and the consequent contribution that the utterance is to make (Stone & Thomason 2002).

The revealed intention is thus a rational reconstruction containing much detail that the speaker may not have consciously attended to. Rational reconstructions of this sort are another idea that, in pragmatics at least, go back to Grice. But they are also common in applications of cognitive architectures to nonconversational reasoning.

Abduction can be implemented in several ways. One way is to take a probabilistic approach, and frame abduction as a form of statistical inference. Indeed, in related work, we have developed a probabilistic version of the

dialogue agent described in this article (DeVault 2008, DeVault & Stone 2009). This version of the agent uses a statistically trained abductive interpretation model to assign probabilities to competing utterance interpretations. The model assigns quantitative costs to alternative assumptions about aspects of the speaker’s intention, where the detailed costs have been tuned by machine learning techniques using a corpus containing thousands of example speaker utterances.

In addition to probabilistic approaches, abduction can also be implemented in a more qualitative way, treating explanation as a form of theorem proving. It is very natural to combine this idea with logic programming methods, which use rules (logical conditionals conforming to certain syntactic restrictions) and data (atomic formulas) to search backwards from a goal to a proof of the goal from the data using the rules.

An abductive theory consists of rules and data, but the rules associate COSTS with their premises. The theorem-proving algorithm of a logic programming method like Prolog is then modified to search through ABDUCTIVE PROOFS—proofs that may involve additional hypotheses taken on as assumptions—to find the least-cost abductive proof of the goal. The cost of a proof will depend not only on the costs associated with the rules, but on the available data, and the scoring method favors UNIFYING EXPLANATIONS—hypotheses that explain several different observations at once.

Stickel (1991) developed an algorithm of this sort, which was used as a reasoning engine by Hobbs and his associates. We have modified Stickel’s algorithm to allow speech acts to manipulate the weights, which provides a mechanism for managing attention in discourse; see Stone & Thomason 2002. But abductive theorem proving can also be used to form plans—the actions involved in a plan are the “hypotheses” that are invoked to “explain” the goal of a plan. This makes it possible to use abduction as a unified architecture for both interpretation and generation in discourse, and to show explicitly how the architecture enables interlocutors to coordinate on interpretations as required by Principle (1). See Stone & Thomason 2003.

However it is implemented, abduction is inference to the best OVERALL explanation, taking all the data into account. This methodology requires a holistic approach to pragmatic reasoning. To recognize a communicative intention from a contextualized utterance, disparate pragmatic phenomena, such as disambiguation, anaphora resolution, presupposition accommodation, and implicature detection, must be considered together to determine which is the most preferred overall interpretation.

Although we envisage a unified model of pragmatic reasoning that applies to the listener as well as to the speaker, and so can explain how the reasoning is coordinated in successful communication, our suggestions can best be explained by concentrating on the listener's process of interpretation.

The listener entertains a set of candidate interpretations in attempting to understand an utterance. Each such interpretation represents an intention, and so characterizes the speaker as committed to certain (private) claims about the conversational situation, and to a certain desired contribution to the conversation. The listener will draw only on public information in working out which candidate interpretation the speaker must have had in mind. But, if coordination is successful, the outcome will be one preferred interpretation that captures the reason for the speaker's utterance. Significantly, this interpretation may highlight speaker beliefs about the conversational situation of which the listener was previously unaware. And it will usually be unproblematic for the listener to accept these background commitments of the speaker, just as it will usually be unproblematic for the listener to accept the speaker's main contribution. Our title, ENLIGHTENED UPDATE, refers to such cases, in which the conversational state is updated to reflect the speaker's recognizable private commitments as well as the explicit contribution of their utterance. This paper considers enlightened update as a possible mechanism behind cases in which the audience recognizes and takes up new information from a speaker's presupposition—cases that have been pretheoretically characterized as accommodation. We argue that enlightened update helps clarify the power and scope of ordinary interpretive reasoning in conversation, while calling attention to a range of circumstances in which interlocutors' reasoning is truly exceptional.

1.5 Plan of the paper

This paper is an exploration of the new notion of a publicly recognizable private commitment. We will argue that this notion brings with it a number of advantages over alternative knowledge-level accounts of presupposition, because it relates presupposition so directly to the requirements of collaboration and coordination in conversation.

- It allows us to square intuitions about grammatical meaning with intuitions about pragmatic attitudes. We can see, for example, how the grammar might require a speaker to commit to certain information,

privately, but in a publicly recognizable way, WITHOUT thereby requiring the speaker to somehow treat it as public, prior information. This gives an attractive way to resolve the well-known and frequent gaps where information must be grammatically backgrounded but need not be shared information among interlocutors. Classic examples include the “informative presuppositions” of change-of-state verbs, factives, and definite noun phrases (Hobbs et al. 1993, Simons 2001, Roberts 2003, Abusch 2010).

- It allows us to narrate pragmatic inference consistently in intuitively appealing ways. For example, we can carry over much of the elegant formal treatment of context change in the presence of accommodation, as developed by authors such as Beaver (2001) and Barker (2002), which is couched in terms of uncertainty about the context.¹

We can now link this account to models of collaboration and the associated reasoning about real-world task domains in ways that explain HOW a listener can come to be uncertain about what a speaker might recognizably commit to in an utterance. This places the account on a firmer footing by connecting the formalism to an independently motivated knowledge-level characterization of the problem of interpretation, and clarifying when interlocutors will understand accommodation as an unmarked strategy and when accommodation will trigger marked interpretive effects such as implicatures or repairs.

- Finally, it allows us to IMPLEMENT straightforward models of pragmatic reasoning that respond to a wide range of theoretical and empirical

¹ There remains, however, room for many positions on how much uncertainty there is about the context in the course of a conversation, and about how mutual attitudes can be supported in the presence of this sort of uncertainty. Our own position can be regarded as a special case of Beaver’s very general framework—a special case that places the following three constraints on allowable uncertainty. (1) In our model, the preferences that rank alternative contexts are shared by the conversants and are systematically related to the (abductive) reasoning that supports generation and interpretation. (2) Hence, we can assume that normally, the speaker will minimize ambiguity by presenting the hearer with a most favored alternative interpretation, and that when the speaker fails to do this a clarification subdialogue is likely to ensue. (3) This enables the further assumption that the dialogue returns regularly to CHECKPOINTS, where mutuality is achieved or at least assumed. These three constraints are reflected in the non-probabilistic version of our dialogue agent, which we present in this article, and are largely carried over into the probabilistic version of the agent as well (DeVault & Stone 2007, DeVault 2008, DeVault & Stone 2009).

considerations. In fact, we were led to the account in part as a natural outgrowth of our ongoing efforts to implement a collaborative approach to language use (Stone & Thomason 2003, DeVault, Kariaeva, Kothari, Oved & Stone 2005, DeVault & Stone 2006, 2007, DeVault 2008, DeVault & Stone 2009). We support our discussions with a description of pragmatic reasoning in a simple implemented conversational agent. We have found that by modeling uncertainty in speakers' possible commitments, while making sure that their actual commitments can be resolved by public information, we have been able to carry over important insights from previous work on collaborative dialogue (Larsson & Traum 2000, Stone, Doran, Webber, Blears & Palmer 2003, Purver 2004) while achieving more flexible and more natural interaction.

The remainder of the paper presents our theory of enlightened update in more detail. Section 2 situates our project, by introducing some received analyses of pragmatic phenomena from Linguistics and the Philosophy of Language, and showing how holistic models of interpretation from Artificial Intelligence offer a new way to reconcile the key intuitions at play. Section 3 motivates the particular approach we take through a fine-grained analogy between enlightened update in silent collaboration on physical tasks and enlightened update in cooperative conversation. Section 4 makes this approach concrete: we sketch the architecture, algorithms, and implemented behavior of a conversational agent that uses enlightened update to interactively identify visual objects with a human user, in a "collaborative reference" task similar to that explored by Clark and colleagues (Clark & Wilkes-Gibbs 1986, Brennan & Clark 1996). Finally, Section 5 builds on the perspective we articulate in Sections 3 and 4 to show how the new distinctions we draw can inform and clarify ongoing debates in the theory of presupposition. In particular, we emphasize how explicit accounts of pragmatic reasoning may force us to conclude that phenomena that exhibit suggestive pretheoretic similarities may nevertheless turn out to derive from different underlying mechanisms.

We are excited about the rapprochement we envision between theoretical analyses and AI models. Hopefully, this combination of forces can not only bring more coherent theories but can leverage theoretical insights more directly in the design and construction of conversational systems.

2 Background

The literature distinguishes two explanatory roles for pragmatic theory, within which we must frame any theory of presupposition.

The first challenge is to characterize the interface between semantics and pragmatics. It is a truism that our knowledge of linguistic meaning underdetermines the content that an utterance will convey on a particular occasion. Pragmatics is what gets us from meaning to content. A good way to conceptualize the connection is that pragmatic reasoning will start from an underspecified meaning containing free parameters and will derive fully specified content by instantiating parameters to suitable values retrieved from the conversational situation. The values cannot be retrieved arbitrarily, because they are in fact generally subject to grammatically specified constraints. Thus, the theory of presupposition must have a place for such constraints, and we survey what we know of such an account in Section 2.1.

The second challenge is to fit pragmatic reasoning into a broader account of communication. In particular, successful communication requires interlocutors to AGREE on the content that is contributed to conversation. Insofar as recognizing the content requires interlocutors to draw on additional background information, the information they actually use must be public. Otherwise, interlocutors cannot be sure they are using the information their partner will expect; miscoordination becomes a possibility. Thus, the theory of presupposition must have a place for the public information that plays this distinctive role in coordination, and we survey the principles of such an account in Section 2.2.

In general, we accept the intuitions and frameworks in the literature which have been developed in response to these two challenges. Except, as we argue in Section 2.3, we reject the common assumption that the same mechanisms are at play in the two cases; we distinguish between the grammatically specified information that LINKS content to meaning in the conversational situation and the public information that interlocutors use to IDENTIFY content from meaning in a coordinated way. A major constraint on our theory, we argue in Section 2.4 (and on any theory that does justice to the evidence about pragmatic reasoning), is that it must model the reasoning as a single, overall process, encompassing everything from “low-level” disambiguation to “high-level” inferences about relationships among utterances in discourse, and embracing both reasoning about language and reasoning about the world. In Section 2.5, we show how the computational model of interpretation as

abduction meets this holistic challenge. This background sets the stage for our subsequent discussion—a mix of analogy, analysis and implementation that, we think, shows why our account is not just a conceivable view but an inevitable consequence of the kind of collaboration interlocutors undertake in conversation.

2.1 Grammatical constraints and meaningfulness

We talk about the world around us, and the grammar allows us to use information in the conversational situation to derive the content of an utterance on a particular occasion. Utterances containing demonstratives, as described by Kaplan (1989), offer an illustrative example. Suppose a speaker, call him Andrew, points at a pot on the stove in which potatoes are simmering, and presents his interlocutor Bess with an utterance of (4).

(4) Those are heirlooms.

He thereby means that **p**, those potatoes, the very objects he has pointed at, are **heirlooms** (namely, of a plant variety developed before 1950 for attributes like hardiness and flavor suitable to historic small-scale pesticide-free farming). Thus, the information that Andrew expresses with this utterance explicitly concerns these individuals from the conversational situation.

To plan and recognize such utterances, language users rely on the same knowledge of language across all situations. Words like *those* have a consistent linguistic meaning that underwrites the content they contribute on different occasions. The content of *those* requires an act of demonstration whose object is a plurality; when this requirement is met, this plurality serves as the semantic value of *those*, and compositional semantic rules determine how this value contributes to the content of the uttered sentence.

For the discussion that follows, it will be useful to indicate how an explicit theory could state such generalizations about (decontextualized) meaning and content. Grossly simplifying and specializing the relevant syntactic and semantic principles, we might spell out the theory required to link utterance

- (4) to its content as in (5).
- (5)
- a. Any utterance of the expression *heirlooms* is an indefinite plural NP and refers to the property **heirlooms**.
 - b. If uttered in association with a demonstration that designates the objects **S**, then the definite plural NP *those* refers to **S**.
 - c. If E_1 is a definite plural NP that refers to **U** and E_2 is an indefinite plural NP that refers to **P**, and E_1 and E_2 are part of a larger utterance of E_1 *are* E_2 , then this larger utterance is a sentence and expresses the proposition that objects **U** have the property **P**.

The theory in (5) has as a logical consequence the semantic characterization of (4) given in (6):

- (6) If uttered in association with a demonstration that designates the objects **S**, then *those are heirlooms* expresses the information that objects **S** have the property **heirlooms**.

Note how the formal structure of (6) represents the link between meaning and content parametrically. The meaning of the utterance includes a free parameter **S**. The grammar constrains the value of **S** based on the associated demonstration and the objects it designates in the conversational situation. We obtain the content of Andrew's utterance from its meaning by instantiating **S** to **p**, taking into account the fact that Andrew's demonstration in fact indicates **p**. Thus we derive the specific content that (4) expresses the information that **p** have the property **heirlooms**.

In framing our intuitions about meaningfulness, it is crucial to sharpen our understanding of the constraints and processes that link parameters and values in such cases. We assume that the reference of the demonstration is DETERMINED by the speaker's intentions. Even if the speaker's act of pointing is poorly aimed, and the hearer is plainly unable to see this act, under normal circumstances the demonstration indicates the potatoes if the speaker intends it to. On this story, your intention in demonstrating an object determines what you refer to with a deictic pronoun in much the same way that your intention to go to your refrigerator determines your destination. Within wide cognitive limits, you can choose any destination you want—but the success of the choice depends in part on the wider world in which the choice is made. (You will have no destination if the refrigerator has been destroyed. You will have a destination if, unknown to you, the refrigerator belongs to someone else—but it will not be your refrigerator.)

Thus, for an utterance of (4) to be MEANINGFUL, it just needs the designated objects to flesh out its underspecified meaning to a specified proposition. The resolution of this underspecification follows from the facts about the communicative action and the situation in which it is performed.

Reasonable intentions are constrained by their place in plans, and in particular by the goals that they serve. A demonstration that does not contribute to the conversational goals may determine a reference, but it is still fundamentally flawed. And even quite reasonable communicative intentions can fail if the circumstances don't cooperate, for example if the audience is ignorant or misinformed about key background. Take (4), and imagine that Andrew has pointed to a covered pot, designating its contents. His interlocutor Bess may not know what's cooking. It might be the potatoes bought at the farmer's market this morning; it might be the beets. One way to describe this is that Bess incompletely identifies what Andrew means. Andrew in fact means that the potatoes are heirlooms. That's what's in the pot. But Bess doesn't know whether Andrew means the potatoes are heirlooms or means the beets are heirlooms. If Andrew really wants to inform Bess that the potatoes are heirlooms, this communicative action will not do it. In this sense, his intention is flawed.

This uncertainty can infect the speaker as well as the hearer. Suppose that Andrew himself uses (4) without knowing whether the pot contains the potatoes or the beets. Again, we can defend the intuitive picture that Andrew still does mean that the potatoes are heirlooms. That's what he's pointed at. Still, he has expressed a proposition that he himself cannot identify. Indeed, in general, it seems that we cannot assign public meaning to utterances—as our intuitions invite us to do—without recognizing cases where speakers and hearers are ignorant of crucial matters of meaning (Kripke 1972, Burge 1973, Putnam 1975). These intuitions are broadly accepted in the literature in the Philosophy of Language. We would like a pragmatic theory that endorses and explains them.

2.2 Coordination

At the same time, we take seriously arguments that mutuality of information is central to pragmatic reasoning. The importance of mutuality is supported by two sorts of theoretical and (loosely) empirical considerations. The theoretical reasons were developed by David Lewis (1969), and apply generally to cases requiring the coordination of actions over a population of rational

agents. Using the theory of cooperative games, Lewis shows that the adoption of an ideally rational, CONVENTIONAL signalling system requires mutual knowledge.² Lewis's insights have become the basis of theories of economic transactions (Aumann 1976, Geanakoplos 1994) and communication (Fagin, Halpern, Moses & Vardi 1995). These theories reinforce Lewis's arguments that rational cooperation is impossible in many cases without mutuality.

Similar theoretical considerations have made their way into computationally inspired theories of discourse (Perrault & Allen 1980, Joshi 1982). Well known reasoning problems like the muddy children puzzle, although artificial in some ways, can be made to seem plausible and compelling, and depend essentially on the assumption that a public announcement to a group creates mutual knowledge among the members of the group.

Also, it is generally accepted that a distinction between given and new information is important in understanding many linguistic phenomena, and there are many examples in the literature seeming to show that mutual knowledge captures the appropriate notion of mutuality in many cases. One particularly convincing and well presented example of this sort is the "Monkey Business" scenario in Clark & Marshall (1981), showing that a felicitous utterance of *Have you seen the movie playing at the Roxy* requires the speaker to believe that the reference of *the movie playing at the Roxy* is mutually known.

Such examples are especially natural and convincing when one takes into consideration only the speaker's beliefs about the hearer's attitudes: you should not ask someone to hand you the crescent wrench unless you expect that they know what a crescent wrench looks like, and can see the wrench. But true mutuality requires arbitrary iterations of attitudes; for instance, *A* and *B* do not mutually know *p* unless *A* knows that *B* knows that *A* knows that *B* knows that *p*. Common-sense reasoning tends to break down over such iterations, and it takes a certain amount of work even to construct examples involving more than two iterations. This makes it difficult to construct a compelling empirical case for the mutuality of given information in discourse. It is the combination of simple examples with the difficulty of finding any principled place to cut off the need for iterations, and with the importance of mutuality in economic theories, that makes mutuality such an attractive concept in the theory of discourse. But the deployment of this assumption

² We use the terms 'knowledge' and 'belief' here, depending on which tends to sound more natural. But the relevant attitude needn't be knowledge or belief. It could be supposition, or any attitude that serves to guide actions in a given, circumscribed situation.

in discourse has to be qualified in the same way that it is in economics; this is an assumption about ideal rationality, that has to be applied with care to human reasoning in realistic situations.

Some authors, such as [Sperber & Wilson \(1995: 15-21\)](#) have claimed that, since a mutual belief has infinitely many logically independent consequences, we cannot have mutual beliefs. If we took such arguments seriously, we would have to conclude that no observations could lead us to conclude that we have ten fingers. Nevertheless, if we require mutuality in discourse, even as an idealizing assumption, we need to show how it is possible to reason plausibly to conclusions about mutuality, and maintain mutuality in the course of a dialogue. We believe that this can be done. For an account of mutuality in terms of default suppositions, see [Thomason 2000](#).

2.3 Reconciling these intuitions

Approaches deriving from Stalnaker's publications on presupposition tend to share the idea that presupposition is primarily attributed to speakers (not to sentences) and that it is related to the common ground, a mutual attitude that tracks what has been established in a conversation. One form of the idea, introduced in [Stalnaker 1975](#) and advocated in some of Stalnaker's later works, is that a speaker's presuppositions are the things she believes belong to the common ground. This idea has the advantage of simplicity, and is workable in highly idealized conversations. But it is very difficult to apply the idea to actual conversations, where misunderstanding or incomplete understanding is a possibility; see [DeVault & Stone 2006](#) for discussion. We need a more robust characterization of presupposition to develop a realistic account of conversation and a detailed account of the reasoning that underpins it.

Instead, we work with a more general notion of presupposition deriving from the structure of intentions, which we take to involve not only a GOAL (a state of affairs to be achieved), but a PLAN (a partially specified way of achieving the goal), and PRECONDITIONS, or ways that the world is assumed to be, and on which the success of the intention depends. Communicative intentions, like other intentions, can have preconditions. For example, as described in [Section 2.1](#), the grammar may encode preconditions about the conversational situation that are required for an utterance to be meaningful. Accordingly, a speaker's communicative intention in using that utterance must take on the commitment that these preconditions are met. Commu-

nicative intentions can also have preconditions that must be met so that a meaningful utterance will have its intended effect. For example, suppose Andrew asks “Did any of your students solve the problem?” and Bess answers “Kim did”. In using this utterance, Bess commits to a specific grammatical analysis: her utterance picks out a specific individual *k*, and contributes that *k* solved the problem. But (on the natural interpretation), Bess’s intention with this utterance goes further. She means to answer Andrew’s question in the affirmative. The information she provides can do this only if the individual *k* is in fact one of her students. Bess’s communicative intention correspondingly incorporates this precondition. It is common for theories of speaker presupposition to treat these two kinds ingredients of utterance interpretation on a par. Stalnaker’s does so, and ours does as well. This is why we see our approach as offering a uniform account of presupposition, even though the linguistic literature tends to focus on the preconditions that are “triggered” by grammatical rules, whereas our formalism and implemented case studies most prominently feature commitments associated with preconditions of discourse coherence.

In sum, we identify preconditions of both kinds as speaker presuppositions. While we acknowledge that speakers will track the common ground to ensure their intentions are recognized—clearly an intention that cannot be recognized is hopelessly flawed—we do not go the further step of assuming that the speaker must take the preconditions they commit to be part of the common ground. On our view, the speaker always has the option to commit to preconditions that offer new information about the conversational situation, as long as the audience can recognize the speaker’s intention. When they do so, they discover what the speaker’s commitments are, and can take up this information. ENLIGHTENED UPDATE is our name for this kind of unproblematic uptake of informative presuppositions.

Analyzing presupposition via intention is similar in some ways to an alternative account of speaker presupposition that Stalnaker mentions (1973), according to which a speaker presupposes the things she is disposed to take for granted.³ This alternative is not as prominent in Stalnaker’s work or as influential as the one that appeals to common ground, but it has the advantage of allowing straightforwardly for informative presupposition. As far as we know, clarifying what it is to “take something for granted” in terms of intention is a new idea. This characterization grounds presupposition

³ Stalnaker’s views on presupposition have evolved over a long period of time, and are complex. For a readable and informative discussion of this topic, see [Simons 2003](#).

in the independently-motivated principle that meanings are intentions of a certain sort. And—most important, for our purposes—it provides insights into the related reasoning processes, since the intention that stands behind an action of any kind will explain the action by providing a goal, a method of achieving that goal, and certain preconditions. This explanation will involve knowledge of the same sort that is needed for planning: knowledge of the effects and preconditions of actions. It is through formalizing this reasoning, and the knowledge it depends on, that we are able to reconcile our view of presupposition with the requirement, explored in Section 2.2, that interlocutors coordinate using mutually supposed information.

Treating informative presupposition, as we do, as a matter of enlightened update, eliminates the need to explain the main body of phenomena either in terms of repair strategies arising from the violation of a rule (Lewis 1979), or in terms of pretense (Clark 1992), or in terms of sequential updates of the common ground in which presuppositions are added before the primary assertion (von Stechow 2008). All of these ideas have some plausibility, and may prove to be useful parts of pragmatics. But the first two of them, we feel, are not appropriate for a phenomenon that is as common and routine as informative presupposition. And we know of no independent way of motivating multi-stage updates directly from considerations having to do with pragmatic reasoning. Without motivation of that kind, it looks very much like an *ad hoc* solution, whose only purpose is to save an account of speaker presupposition that requires presuppositions to belong to the common ground.

In Stalnaker 2002, Stalnaker suggests a view of accommodation as a side-effect of conversants' attempts to cooperatively align their presuppositions in the course of a conversation. This is very close to the ideas we discuss in Section 4.4, but again, we are able to relate this to an account of the reasoning that is much more detailed.

In fact, our insistence on a model of the reasoning that is sufficiently detailed to support an implementation is perhaps the most important difference between our approach and others in the literature, since it shapes our position in fundamental ways. This methodology requires us to choose a reasoning domain, and to deal systematically with the cases of presupposition accommodation, as well as the other pragmatic phenomena, that arise there. Since tractable reasoning domains tend to be relatively simple, though they provide many examples of presupposition accommodation, these examples will seem relatively simple and routine. We are betting that a

methodology working with such examples, but accounting for the reasoning with unprecedented detail, will provide a valuable counterpoint to more traditional linguistic methodology.

2.4 Pragmatics and cognitive modularity

Linguistic methodology favors an approach that leads from intuitions about linguistic examples to structural generalizations, which then are explained in terms that are (1) modular and (2) independent of the reasoning mechanisms. Thus, a grammar of a language is a knowledge-level account of the language's structures. The grammar is supported by reasoning procedures—for instance, the procedures involved in parsing—but there is no reason to expect these procedures to be similar to and continuous with the procedures that are involved in general-purpose common-sense reasoning. Linguists do not in general seek to integrate their theories with knowledge-level accounts of nonlinguistic domains, and usually modularize their grammars into more or less independent subsystems.

These practices are not well adapted to pragmatics. Task-based discourse is closely connected to the domain reasoning associated with the task. Pragmatic phenomena such as reference resolution, disambiguation, implicature detection, and the identification of discourse relations depend heavily on domain knowledge. And pragmatic reasoning is often similar or even identical to domain reasoning. Suppose Bess asks Andy, who is standing by, for the salt. If Bess is eating, Andy correctly infers that Bess wants the small salt shaker. If Bess is cooking and holding a tablespoon measure, Andy correctly infers that Bess wants the box of salt. Andy's reasoning here is identical to or similar to the reasoning he employs in planning how to use salt for various purposes. Examples such as this are commonplace.

We can't hope, then, to isolate pragmatic reasoning from domain reasoning. Neither can we assume that pragmatic reasoning divides into disconnected special purpose modules, devoted to different pragmatic phenomena, because the preferences that govern the selection of a favored interpretation are OVERALL preferences. The following discourse illustrates the point.

- (7) Ann thought she might be able to get help with her computer problems from Betty. She is a systems analyst.

The fact that *Ann* is the subject of the first sentence creates a *prima facie* preference for Ann as the reference of *she* in the second sentence. However,

the two sentences can be regarded as an explained/explanation pair if *she* is resolved to *Betty*, and a plausible implicature—that Ann thought that a systems analyst would be helpful with her problems—would also be generated. These factors create an overall preference for Betty as the reference of *she*. Many similar examples can be found in [Hobbs et al. 1993](#).

Thus, the reasoning associated with various pragmatic domains (anaphora, reference resolution, metonymy resolution, implicature, presupposition, etc.) must produce commensurable preferences, and must access and overlap with domain reasoning. Under these circumstances, we believe that the best working hypothesis is to account for the reasoning in as non-modular a way as possible: pragmatic reasoning uses knowledge structures and reasoning that are similar to those used in domain reasoning, and is comprehensive, dealing at once with all pragmatic phenomena.

In [Grice 1957](#), Paul Grice assimilated discourse meaning in cases involving language with nonlinguistic cases involving only cooperation and intention recognition. Endorsing and adopting this theme, we believe that pragmatic reasoning will have much in common with the reasoning that supports collaborative human activity.

2.5 Interpretation as abductive intention recognition

Abduction provides a reasoning framework for pragmatics that is properly general-purpose, and that supports domain reasoning such as explanation and planning. The knowledge structures are intentions and their ingredients, and preferences over these intentions. The framework inherits the naturalness and plausibility of intention-based approaches to pragmatics, and with some natural extensions concerning helpfulness and the maintenance of mutual information, it can also be used to support collaboration.

The logic programming approach developed by Hobbs and his associates to implement this knowledge-level picture of interpretation as abductive explanation of the utterance provides an equally appealing account of the reasoning. Hobbs took proving the logical form of the uttered sentence to be the goal of the hearer's reasoning,⁴ in effect taking the hearer to be asking the question "What does this sentence mean?" We prefer to think of the goal as proving that a certain intention supports the uttered sentence; in effect, the hearer is always asking "What did the speaker intend by saying that?"

⁴ Hobbs used a somewhat idiosyncratic version of logical form, with some similarities to linguistic LF.

As we said, an abductive proof will in general require new assumptions, which are added to the conversational agent's stock of working assumptions. This provides a way of distinguishing the new information in an utterance (the assumptions created by the process of interpretation) from old information (the stock of working suppositions that are used as background for the interpretation). The distinction is similar to the one deriving from Stalnaker between the pragmatic presuppositions of an utterance and its new information, and it is formulated in an implementable reasoning framework that combines grammatical requirements for wellformedness and meaningfulness with the informational demands of domain reasoning and conversational coordination. Thus, Hobbs's abductive approach shows how the grammatical requirements for meaningfulness and informational demands of coordination can be integrated in conversational reasoning.

Thinking of interpretation as the derivation of an INTENTION rather than of a sentence meaning introduces a more subtle set of distinctions into pragmatic theory. An intention can introduce propositional assumptions in various ways. If, for instance, the intention is assertional, it will have a primary goal of updating the common ground with a proposition. But it may also have preconditions—propositions that are required for the feasibility or appropriateness of the intention—which can be new information, even though they are not part of the content of the assertion.

The present paper, like our earlier ones (Stone & Thomason 2002, 2003), is an attempt to put these different ingredients together.

3 Collaboration and conversation

As we explained in Section 2.4, we subscribe to the methodological principle that pragmatic reasoning is continuous with domain reasoning. We now expand on this by developing the idea that reasoning about conversational intentions parallels reasoning about intentions in other kinds of collaborative activity. In fact, the collaborations that we undertake as embodied agents in the physical world, because they lack the abstraction of linguistic interaction, often offer the best illustration of pragmatic principles, and of the role of abductive intention recognition in the reasoning.

We wish to emphasize the following points concerning intentional collaborative activity:

- (8)
 - a. Any intentional activity, collaborative or not, requires us to make private commitments about how the world is.
 - b. Coordinated activity requires keeping track of mutual information, and distinguishing it from what is private.
 - c. Mutual information must be updated and maintained, along with private beliefs.
 - d. Agents capable of intention recognition can transmit information by performing actions in public.
 - e. If the mechanisms of intention recognition are mutual, public actions can create new mutual information.
 - f. In particular, an agent's public performance of an action that is mutually known to require a commitment C for its successful performance will add to the mutual information the proposition that the agent believes C . Normally, this will result in the addition of C to the stock of mutual information.

We illustrate these points with an extended example. Suppose Andrew and Bess are cooking. Bess asks Andrew to wash and dry some spinach for a salad. Bess then turns her back to Andrew and attends to a pot on the stove. It is boiling ferociously, and she can neither see nor hear what Andrew does. Andrew cleans the spinach. Bess turns to face Andrew again.

In this situation, we have to distinguish between the state of the collaborative activity, the agents' private knowledge, and their mutual information. In the current state, Andrew has cleaned the spinach and knows that he has done so. But Bess does not know whether Andrew has cleaned the spinach. And we can suppose that Bess and Andrew both mutually recognize that Bess does not know this. So it is mutually recognized at this point that maybe Andrew has cleaned the spinach but maybe not.

Let us continue the example. In plain view, Andrew mixes the spinach with the other salad ingredients. In fact, this action furthers the joint activity in a straightforward way; Andrew and Bess are one step closer to a good salad.

Consider how Bess updates her information as she watches this. For all she knows, Andrew may not yet have washed the spinach. But on seeing Andrew add the spinach to the salad, Bess can rule out this possibility. Andrew himself is in a position to know whether he has washed the spinach, and is committed to making a good salad. So if he hadn't washed the spinach, he wouldn't have added it to the salad. Andrew's action thus

eliminates the ambiguity in the state of the activity. Afterwards, because of this ENLIGHTENED UPDATE, Andrew and Bess enjoy the mutual information that the spinach is washed and clean, and so the meal preparation continues smoothly with this as common ground.

We can compare this picture with *Stalnaker's* (1978) characterization of assertional acts as speech acts that add a proposition to the common ground. In our framework, assertion is not the only method of incrementing the common ground: the mechanism described in Characteristic (6) provides another way, based on intention recognition, for intelligent collaborating agents to add private commitments that can be mutually inferred from shared information (including mutually observed actions) to the common informational background. Indeed, updates of this kind do not even require verbal communication.⁵ However, we can supplement this mechanism with explicit verbal communication; we can and sometimes must provide verbal information to accompany what we do, so that our private commitments are correctly recognized.

Andrew's action of mixing the salad was public, and its performance was mutual to the participants. While Andrew's action of washing the spinach was hidden, but Bess's inference about that action used only mutual information about salad making. So Bess can infer that the conclusions she draws about the action belong not only to her private store of information, but to the mutual common ground.

In this setting, Andrew's action of cleaning the spinach is what we call a TACIT ACTION. An action in a collaboration is tacit if some of the collaborators do not immediately know whether it has taken place or whether it has had the appropriate effects. We will call other actions PUBLIC ACTIONS; collaborators know when public actions occur and whether they succeed.

Tacit actions give us clear grounds to distinguish between the state of the activity and participants' information—they create a gap between what participants have done and what they mutually know. As we will show, in our computational model, a speaker's commitment to the fact that some tacit action has occurred is a common type of PUBLICLY RECOGNIZABLE PRIVATE COMMITMENT; these commitments are identified through the agent's abductive intention recognition process.

⁵ Nor are they meant, in the sense of *Grice* (1957). In this cooking example, Andrew does not mean by mixing the spinach in the salad that he has washed it; he does not intend Bess to recognize that he has washed the spinach in part because of her recognition of his intention for her to recognize this.

Now, we do not claim that all such gaps between mutual and private knowledge about an ongoing activity are due to tacit actions. There may be exogenous events that some collaborators can observe, but some cannot. For example, perhaps only Bess will see when the contents of her boiling pot have cooked. This would be a tacit event, not a tacit action. Andrew might still infer that this event has taken place from observing Bess's actions—say, by watching Bess turn off the heat or empty the pot. Moreover, there may be uncertainties about the state of the collaboration that aren't usefully traced back to any particular event in the collaboration at all. For example, perhaps it's Bess's kitchen and so only she knows that the small ceramic bowl on the countertop is a salt cellar. Andrew might discover this by observing Bess adding a pinch of white crystals from the bowl to the pot on the stove. Thus, Andrew's reasoning in explaining Bess's actions may very well be similar whether he is resolving uncertainty that arises from tacit actions in the collaboration, from unobserved exogenous events, or simply from incomplete background knowledge. Nevertheless, we focus on the case of tacit actions in this paper. They are a clear case, they are an indispensable ingredient of problem-solving interactions, and they have, we think, a particular significance for conversation.

In this paper, we explore the agents' reasoning in response to a mix of public and tacit action in a collaboration. When a public action follows some tacit actions, interlocutors can fill in gaps in their knowledge by exploiting their understanding of the ongoing collaboration. Whenever a public action is taken, the collaborators all discover that the state of the activity already made the public action appropriate to the state in which that action was taken. If the public action were not appropriate without certain changes to the environment, its agent should not have so acted. If the hearer has not observed these actions, they must be among the tacit actions that preceded the public actions. In many cases, like our example, it is obvious what these tacit actions must have been. The collaborators can thus jointly infer that the tacit actions must have occurred and, in relevant respects, succeeded. As a result, collaborators can make an enlightened update to their information state—they can proceed in effect as though they always knew about the tacit actions and their relevant effects.

Tacit actions, and the distinction between mutual and private information, are regularly taken into account in reasoning about collaboration and conversation. In particular, mutual beliefs are crucial in effective collaboration, and as Herbert Clark and others have noted, speakers are careful to maintain

mutuality. Suppose that at this point in the cooking activity, it is mutually known that a pinch of cayenne and a tablespoon of cumin need to be added to the pot. The spice containers are not labeled, but Bess knows which is which. She hands a container to Andrew. Given mutual information about the task, Andrew can infer that he is to add the spice to the pot. But Bess's previous contribution was flawed, since at this point Andrew doesn't know which spice it is and so doesn't know how much to add. By keeping track of shared information, Bess can infer that simply handing the spice to Andrew would be uncooperative. Repairing this flaw could lead to disambiguation by adding a verbal clue—say, by saying “here's the cumin” as she passes the spice container.

Of course, people engaged in cooperative activity can make mistakes, and produce ambiguous and even misleading actions, just as they sometimes produce ambiguous and even misleading utterances. But the fact that we classify these cases as mistakes shows that we recognize the importance of tracking mutual information in these cases.

If Bess is a worrisome type, or feels that Andrew is less than conscientious about washing vegetables, she may ask Andrew for explicit confirmation that he indeed washed the spinach. But if this is a task they have performed many times, and Bess has no worries about Andrew's competence, an enlightened update can add the occurrence of a tacit action to the common ground, just as if the action had been verbally confirmed or publicly observed. Bess can say things like

(9) Did you wash the parsley when you washed the spinach?

for instance, with perfect propriety. Our framework thus explains one way a speaker can anticipate that a specific presupposition, like *you washed the spinach*, will be straightforwardly accommodated by an enlightened update. In collaborative activity, enlightened update not only serves as a valid way to add information to the common ground, but it is frequent and pervasive.

4 Pragmatic reasoning in a cooperative planning domain

In this section, we flesh out our case by presenting a simple model of the coordinated reasoning of collaborators in conversation. As in our previous papers (Stone & Thomason 2003, Stone 2004b), we consider two-party interactions, and conceptualize the generation process of the one agent *S* and the understanding process of the other agent *H* in symmetrical ways. The

reasoning problem that grows out of the discussion of Section 3 is how to handle the uncertainty about the state of the ongoing task that arises because of the collaborators' incomplete knowledge about the actions that have been performed.⁶ In particular, we must explicate both the intentions the two collaborators can commit to, and the evidence they obtain from observing each other, in ways that are sensitive to the different information that each each brings to the ongoing task. Our understanding of presupposition as a component of speakers' communicative intentions ties these reasoning models back to the pragmatics of presupposition and accommodation.

Concurrently, we describe how we have implemented this model in a conversational agent, COREF, that works together with its interlocutor to identify objects from a scene displayed in a graphical computer interface. Our demonstration system plays a referential communication game, much like the one that pairs of human subjects play in the experiments of [Clark & Wilkes-Gibbs \(1986\)](#). Two subjects are presented with a display containing the same set of objects. One, the DIRECTOR, sees them arranged in a specific linear sequence, while the other, the MATCHER, sees them shuffled arbitrarily across their visual space. The interlocutors must go through the objects one by one so that the matcher lays them out in the director's order.

We describe each episode in this game as an activity involving the coordinated action of the two participants. The director *D* knows the referent *R* of a target variable *T* and the matcher *M* needs to identify *R*. COREF can play either role, *D* or *M*, using the objects in the graphical display as candidate targets and distractors, and using text as its input and output.

We exemplify what COREF does in (10). The corresponding graphical display is shown in Figure 1.

- (10) C₁:This one is a square.
U₂:Um-hm.
C₃:It's light brown.
U₄:You mean like tan?
C₅:Yeah.
C₆:It's solid.
U₇:*[Moves the object into position]*.

Here COREF (C) and the user (U) exchange six utterances—and one action—in the course of identifying the tan solid square that must fill the next square.

⁶ The thought experiments of Section 3 suggest that language users may face other sources of uncertainty in everyday face-to-face interaction; we leave this for future research.

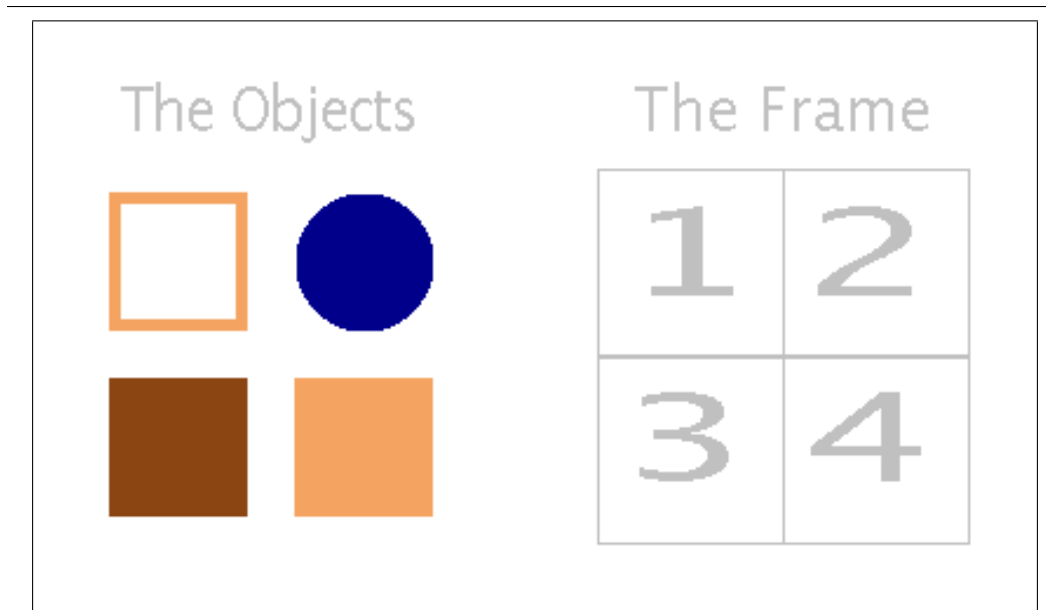


Figure 1 The matcher's view of a collaborative reference task. Four objects are displayed in an unordered array at left. These objects need to be placed, as specified by the director, into the frame at right.

COREF uses the same task knowledge and the same grammar whichever role it plays. It reasons from a specification of collaborative reference that includes tacit actions, and thus makes frequent recourse to enlightened updates. These updates figure not only in utterance understanding and generation, but also in planning and recognizing nonlinguistic public actions, such as the user-interface actions represented by U_7 . Of course, COREF also draws on private knowledge to decide how best to carry out its role; for now it describes objects using the domain-specific iteration proposed by Dale & Reiter (1995). The knowledge we have formalized is targeted to a proof-of-concept implementation, but we see no methodological obstacle in adding to the system's resources.

4.1 Formalizing context in collaborative activity

We model a collaborative planning activity as a dynamical system in which the states of a domain are altered by successive actions. Each point in the activity is characterized by a state s of the domain, and by the knowledge

the agents have of the domain state and of each others' knowledge. The epistemic aspects of such systems can be modeled using possible-worlds semantics; see [Fagin et al. 1995](#). In [Reiter 2001](#), Raymond Reiter shows how to incorporate these ideas in settings of the sort used in Artificial Intelligence to formalize planning.

To model collaborative conversation, in particular, we take each state to specify an abstract, objective body of information that constitutes the conversational situation—the information state (IS) of the conversation ([Poesio & Traum 1997](#), [Larsson & Traum 2000](#)). Thus, it will spell out the attributes of dialogue state that are familiar from computational theories of discourse, including: sets of propositions contributed to the conversational record ([Stalnaker 1978](#)), plans and problem-solving activities that are underway ([Lochbaum 1998](#)), outstanding interlocutor obligations ([Traum & Allen 1994](#)), linguistic forms of prior utterances ([Purver 2004](#)), etc.

Precisely what information goes into the information state depends on what interlocutors are doing. For a task of collaborative reference, the IS tracks how interlocutors together set up and solve a constraint-satisfaction problem to identify a target object. In any point in this problem-solving, *D* and *M* have agreed on a target variable *T*, a set of candidate referents *C* that could serve as the value of *T* and a set of constraints that the value of *T* must satisfy. We call this body of information a CONSTRAINT NETWORK ([Mackworth 1987](#)). When *M* recognizes that these constraints identify the target referent *R*, the task ends successfully. Until then, *D* can take actions that contribute new constraints on *R*. Importantly, what *D* says adds to what is already agreed about *R*, so that the identification of *R* can be accomplished across multiple sentences with heterogeneous syntactic structure.

In COREF's domain, the overall game is to identify a set of objects in turn. Part of the information state—the task state—tracks this activity in terms of episodes of collaborative reference. We represent the task state through a tuple $c^* = \langle R, P, T, C, U, O, H \rangle$, where *R* is a set of objects yet to be identified, *P* is a set of propositions that have previously been added to the context (e.g. by a player action or utterance), *T* is a stack of active tasks (where each task specifies a joint goal that the collaborators are engaged in and the actions they can take in support of that goal), *C* is a set of constraint networks (one for each target object), *U* is the overall universe of discourse (a set of properties and objects), *O* is the set of visual objects present in the interface, and *H* is a historic record of previous events in the dialogue. In this state, *U* depends directly on the visual display; *O* contains all the objects on

the display; R initially is equal to O , but is narrowed down as these objects are identified one by one. P straightforwardly accumulates the information that interlocutors have exchanged. Meanwhile, T and C evolve in a richly structured way according to interlocutors' actions.

This model contrasts with traditional plan-based models, as exemplified by Heeman & Hirst's (1995) model of goals and beliefs in collaborative reference. Our representations are simpler than Heeman & Hirst's but support more flexible dialogue. For example, their approach to (10) would have interlocutors coordinating on goals and beliefs about a syntactic representation for *the tan solid square*; for us, this description and the interlocutors' commitment to it are abstract summaries of the underlying collaborative activity. We formalize the activity as a cumulative and open-ended collaborative process, much as advocated by Rich, Sidner & Lesh (2001).

Our information state tracks other features of the conversational interaction in addition to domain activity. For example, our IS also tracks the salience of entity and property referents and explicitly stores the previous utterance, enabling subsequent reference to it. Meanwhile, the IS also allows subtasks of questioning or clarification that interlocutors can use to maintain alignment. In fact, the same constraint-satisfaction model is used not only for identifying displayed objects as part of the collaborative reference game but also for identifying abstract entities, such as properties, in pursuing clarification subdialogues. The key antecedent for these extensions is Purver's (2004) characterization of clarification of names for objects and properties. We extend Purver's work to allow a cumulative and open-ended treatment of clarification; when we describe things, our descriptions grow incrementally and can specify as much detail as needed.

4.2 Action, tacit action, uncertainty and coordination

The collaborative context evolves over the course of the dialogue through the domain-dependent set of action types, \mathcal{A} , that interlocutors can take. Each action a that interlocutors take has the form $a = \sigma(\alpha)$ where $\alpha \in \mathcal{A}$ is an action type and σ instantiates the free parameters of α . Doing a effects a *deterministic* transformation on the current context. We can formalize this by way of an update function that recalls the notation of the situation calculus:

$$(11) \quad s_{t+1} = \text{do}(a, s_t)$$

(Again, see Reiter 2001.) Thus, in the deterministic case, the state of the domain can be characterized either by specifying the state itself or by specifying the sequence of actions that led to it. Such a sequence will produce a unique state from the initial state of the domain.

In our formalization, we assume that each action type is either PUBLIC, which means that any time any action of this type is performed, its occurrence becomes mutual knowledge among all agents, or TACIT, which means that only the agent that performs an action of this type knows about its occurrence. In other words, our model embodies the simplification that observability is just a matter of the type of action performed, idealizing away from its dependence on the contingent relationships among the agents.

As agents perform tacit actions, they must coordinate with one another to maintain mutual knowledge of the domain state. With COREF, we develop a precise account of one way this coordination might proceed by regimenting how the collaboration takes place. In particular, we imagine that the collaboration intersperses stages of uncertainty with CHECKPOINTS—states in which both agents have full, mutual knowledge of the domain state. We assume that agents achieve this by acting in conformity to a Principle of Coordination Maintenance, stated formally as (18) below. In brief, according to this principle, agents will limit what they do and elaborate how they present it so that they can keep in synch with one another.

Suppose that in the initial state of the entire system, both agents have mutual knowledge of the domain state. At this point one agent, S , performs a sequence of tacit actions. H cannot observe these actions, but can assume that S is acting in pursuit of their joint collaboration. In the subsequent state of the system (1) the domain is in the state that results from performing these actions, (2) S knows that the domain is in this state, and knows that H considers any legitimate sequence of actions performable by S on the initial domain state to be possible, and (3) H knows that S knows the state of the domain, but as far as the domain is concerned knows only that it has reached a state that can be obtained by a legitimate sequence of S 's actions. This space of possibilities can be represented as the set of branches over a tree whose nodes are actions. Under our assumptions, both S and H have access to this space. However, whereas S knows what point in this space they actually occupy, H lacks this information. We call this space THE HORIZON; if c indicates the checkpoint state, $Z(c)$ indicates the corresponding horizon.

After performing a sequence of tacit actions, S chooses an observable action e . This action will advance the joint activity, but crucially, it also has

to coordinate S and H by resulting in a checkpoint.⁷

In understanding e , H 's task is to identify the transition that e effects: what the state of the activity was beforehand, what e itself does, and what the state of the activity is afterwards. This amounts to recognizing the sequence of tacit actions preceding e , and how e changes this context, which can be done by understanding S 's intention in undertaking e , including the private commitments of this intention. S 's task in generation is to produce an action that enables H to perform this recognition. Since H is using abductive reasoning, this means that S must produce an observable action whose preferred explanation is the intention that actually underlies this action.

In COREF, the action set \mathcal{A} has observable task actions including acknowledging a prior contribution (notated by the symbol `ack`); adding a constraint C to the constraint network for a target t (notated `addcr(t, C)`); placing a specified object x at the next point in the desired sequence (notated `place(x)`); inquiring whether some other action a can now be taken (notated `ynq(a)`); and rejecting an action a as impossible (notated `no(a)`). In addition, much of the flow of the activity is modeled through tacit domain actions, including actions of selecting the referent sequence, initiating collaborative reference to a particular target referent t (notated `Push(CollabRef(t))`), marking the target t as identified (notated `Pop(CollabRef(t))`), triggering a clarification subtask by marking an utterance u as incompletely understood (notated `Downdate(u)` following Purver 2004), or reinstating the effect of an utterance once its meaning has been agreed through clarification (notated `Reinstate(u)`). These moves allow COREF to handle clarification, grounding, and other kinds of implicit task progress with a flexible, declarative model and consistent processes of interpretation and generation, based on enlightened update.

In (10) for example, the system represents its moves as successively constraining the shape, color and pattern of the target object. The user's clarification request at U_4 of (10) marks the agent's description of color as problematic and so triggers a nested instance of the collaborative reference task. At C_5 of (10) the system adds the user's proposed constraint. With

⁷ The version of COREF described in this article requires a checkpoint after each utterance, but we have relaxed this requirement in the probabilistic version of COREF (DeVault 2008, DeVault & Stone 2009). Instead, the probabilistic version of COREF is able to tolerate limited amounts of uncertainty across multi-utterance subdialogues. See DeVault 2008 for the full details on the strategies and representations this version of COREF uses to manage its uncertainty about the dialogue state, and an empirical evaluation of how successful its strategies for managing its uncertainty are.

C_6 , the system tacitly assumes the nested clarification task is solved and that the color constraint is reinstated, and then continues the main task by identifying the pattern of the target object. The user's action of moving the object into position at U_7 shows that they have tacitly identified the object and are continuing on with the task.

The full analysis of task actions in dialogue (10) is presented in Figure 2. To take one case, in interpreting the user's utterance *You mean like tan?* (U_4 of (10)), COREF interprets the user's intention as including the following sequence of actions:

- (12) 1. Downdate the effects of the problematic utterance C_3 .
2. Initiate clarification.
3. Start collaborative reference, targeting COREF's intended property S_1 associated with *light brown*.
4. Start a question-answer task in which COREF is obligated to answer a user question.
5. Ask the question:
`ynq(addcr(S_1 , $S_1 = \text{sandybrown}$)).`

Only the last two of these can be seen as directly associated with the semantics of the utterance. The others are tacit actions that are recognized abductively as part of COREF's recognition of a communicative intention that makes sense of the user's utterance. In other words, our formal model of the collaborative activity involved in clarification says that a speaker, in using *You mean like tan?*, presupposes that the previous utterance has not succeeded in its usual effect on the conversation, and that a process of repair of that utterance is underway. These presuppositions are preconditions of the discourse coherence of the clarification question, and are part of the user's communicative intention. They are satisfied in the current state of the conversation, because of the tacit actions through which the user has implicitly adjusted it. Before the utterance, COREF is not aware that these adjustments have been made. But COREF identifies them as publicly recognizable private commitments of the user based on the utterance the user makes. COREF then performs an enlightened update that incorporates the effects of these tacit actions into its conversational state, in effect accommodating the informative presuppositions associated with the discourse coherence of the user's contribution.

COREF uses its action representation not just for domain problem-solving but also to manage aspects of dialogue such as the attentional state. The

Utterance	Task Moves	Resulting Task State	
		Task Stack	Facts
		ClarkGame	square(a) solid(a) sandybrown(a)
<i>C₁: This one is a square.</i>	Push(Collabref(t_1)) (tacit) addcr(t_1 , square(t_1))	CollabRef(t_1) ClarkGame	square(a) solid(a) sandybrown(a) square(t_1)
<i>U₂: Um-hm.</i>	ack	CollabRef(t_1) ClarkGame	square(a) solid(a) sandybrown(a) square(t_1)
<i>C₃: It's light brown.</i>	addcr(t_1 , sandybrown(t_1))	CollabRef(t_1) ClarkGame	square(a) solid(a) sandybrown(a) square(t_1) sandybrown(t_1)
<i>U₄: You mean like tan?</i>	Downdate(C_3) (tacit) Push(Clarify(C_3, S_1)) (tacit) Push(CollabRef(S_1)) (tacit) Push(YNQ) (tacit) ynq(addcr(S_1, S_1 =sandybrown))	YNQ CollabRef(S_1) Clarify(C_3, S_1) CollabRef(t_1) ClarkGame	square(a) solid(a) sandybrown(a) square(t_1)
<i>C₅: Yeah.</i>	addcr(S_1, S_1 =sandybrown)	CollabRef(S_1) Clarify(C_3, S_1) CollabRef(t_1) ClarkGame	square(a) solid(a) sandybrown(a) square(t_1) $S = sandybrown$
<i>C₆: It's solid.</i>	Pop(CollabRef(S_1)) (tacit) Pop(Clarify(C_3, S_1)) (tacit) Reinstate(C_3) (tacit) addcr(t_1 , solid(t_1))	CollabRef(t_1) ClarkGame	square(a) solid(a) sandybrown(a) square(t_1) $S = sandybrown$ sandybrown(t_1) solid(t_1)
<i>U₇: [Moves the object into position]</i>	Pop(CollabRef(t_1)) (tacit) place(a)	ClarkGame	square(a) solid(a) sandybrown(a) square(t_1) $S = sandybrown$ sandybrown(t_1) solid(t_1)

Figure 2 A trace of the task actions and the task states through which COREF represents and participates in the interaction of (10).

grammar, for example, specifies that each subject noun phrase updates the context to give its referent the status in the givenness hierarchy of being IN FOCUS (Gundel, Hedberg & Zacharski 1993). This update makes the referent available as a preferred referent for subsequent pronouns. In Stone & Thomason 2003, we explain in detail how we represent such transitions in the framework of interpretation as abduction.

Many simplifying assumptions are incorporated in this model. The domain representation employs a highly specific formalization of the communication topic, and of the content that needs to be communicated at any given point. In this respect, our domain is much more regimented than ordinary, informal conversation. We also assume that actions are asynchronous, but follow the rather rigid protocol described above.

In fact, the clarification protocol exemplified here, in which clarified utterances are downdated and then reinstated after successful clarification questions, has undergone some modifications in related work on this agent. In the probabilistic version of COREF (DeVault 2008, DeVault & Stone 2009), for an ambiguous user utterance, rather than modeling a clarification question by COREF as beginning with a downdate of the clarified utterance, the agent instead becomes uncertain about which of several interpretations was intended, keeps track of the alternative interpretations, and uses a probabilistic framework to propagate its uncertainty forward as the clarification subdialogue progresses.


Of course, additional details of the computational model could potentially be altered or improved. We are, in effect, following the methodology of Artificial Intelligence, in investigating a complex reasoning problem by investigating how it plays out in simplified, formalizable special cases. Despite the simplifications, we believe that the model is realistic enough to shed light on some of the important problems of pragmatics, and that it can be extended incrementally to cover more complex cases.

4.3 Formalizing interpretation

To represent interpretations we must formalize the complex package of structured information that figures in the speaker's communicative intention. Interpretations must specify the action taken, its description under the grammar, the preconditions that the action depends on and the outcome that the action is to achieve.

We formalize the action using a tree-adjoining grammar (TAG) (Joshi,

Levy & Takahashi 1975, Schabes 1990), where any utterance is derived as the composition of elementary tree fragments that specify a lexical item and all of its syntactic arguments. For example, a schematic TAG derivation of *it's light brown* is as in (13).



We spell out the interface between structure and interpretation for structures such as (13) by formalizing the conceptual apparatus developed in Section 2.1. The grammar may associate any elementary tree with parametric actions that take their effects in the context whenever a speaker uses an utterance containing that elementary tree. The parameters in the action are subject to constraints that must be met in context. In addition, compositional semantics—here implemented via unification, following Gardent & Kallmeyer (2003)—may require trees to share parameters, so that one tree may function to specify the value of a free parameter in another.

The constraints associated with (13) are given in (14).

$$\begin{array}{l}
 (14) \quad \textit{predication}(M) \wedge \textit{brown}(C) \wedge \textit{light}(C) \\
 \quad \quad \textit{in-focus}(X)
 \end{array}$$

These constraints represent instructions to fill the three parameters, M , C and X , through which the update for this utterance will be constructed. For M we must find a relevant task action M that interlocutors normally achieve with a declarative sentence such as this; for C we must find a relevant color that can be described as *brown* and *light*; X is the object the utterance describes, and since the utterance uses *it* to refer to it, the grammar requires that *in-focus*(X) holds.

Note that these constraints link utterances to the contributions that are possible in a specific task. As in DeVault, Rich & Sidner 2004, a KNOWLEDGE INTERFACE mediates between domain-general meanings and the domain-specific ontology supported in a particular application. This allows us to build interpretations using domain-specific representations for referents, for task moves, and for the domain properties that characterize referents. In fact, our architecture allows utterances to trigger an open-ended range of domain-specific updates. We do not factor updates to the IS through an abstract taxonomy of speech acts—instead, utterances directly make domain moves, such as adding a constraint.

In (15), we schematize the actions effected by the utterance:

$$(15) \quad M(X, C(X)) \\ \text{PutInFocus}(X)$$

In other words, we know from the form of the utterance that it will make a contextually relevant predication M that says that referent X has property C . Moreover, since referent X is realized as the subject, the utterance will ensure that X has (and retains) the pragmatic status of being in focus.

Together, (13), (14), and (15) constitute a description of the utterance *it's light brown* as an action: the description says what the action is, what preconditions it has, and what outcomes it achieves. These clauses thus constitute the core of a communicative intention.

To completely spell out a reason to use this action, however, we need to explain how exactly the preconditions of this action are met, and what outcome the action is intended to achieve. In this domain, we can link up (14) with specific instances in the information state of the conversation.

$$(16) \quad \text{predication}(\text{addcr}) \wedge \\ \text{brown}(\text{sandybrown}) \wedge \\ \text{light}(\text{sandybrown}) \\ \text{in-focus}(t_1)$$

These constraint instances simply say the following: first, that adding a constraint is the kind of action that in collaborative reference is achieved through a declarative statement; second, that in this context, this specific color, (RGB value F4-A4-60, or XHTML standard “sandy brown”), can be described as *light* and *brown*; and finally, that at this point in the conversation the target t_1 of the ongoing collaborative reference interaction is in focus. If the speaker commits to this specific link between the utterance and the context, the speaker must also commit to bringing about the specific changes to the domain state in (17).

$$(17) \quad \text{addcr}(t_1, \text{sandybrown}(t_1)) \\ \text{PutInFocus}(t_1)$$

In other words, the speaker must in fact add the constraint that the target of the current collaborative reference interaction is that specific color, and must keep that target in focus.

A system needs a complex structured representation—capturing the information in (13), (14), (15), (16), and (17)—to capture the content of a

communicative intention. In Stone & Thomason 2002, 2003, we show that abductive proofs do in fact constitute such complex representations.

To understand an utterance, we reason abductively to find a way of analyzing the utterance according to the grammar and recognize why (the speaker might think) the corresponding preconditions were satisfied. This gives us a proof that could represent a plan that could possibly be carried out. Before we use this plan to represent the speaker's intention, we must also check that the outcome the plan envisages is in fact a contextually appropriate contribution to the ongoing task.

The possibility of tacit actions highlights the importance of our abductive characterization of interpretation. Consider understanding COREF's utterance *It's solid* as C_6 in (10). In fact, COREF intends this utterance to have an analysis that is broadly analogous in form and content to the analysis of *It's light brown* that we've just seen. That is, the utterance gives more information about the domain object that serves as the value for target t_1 by indicating that the object is drawn with a fill rather than in outline. However, this interpretation is possible only if we ASSUME, abductively, that COREF has tacitly finished the clarification subdialogue that its last utterance contributed to. Otherwise, the appropriate action for the task would be to continue placing constraints to identify the color described as *light brown*.

As always, of course, we need to rule out this alternative interpretation as well as permitting the correct interpretation. COREF's implementation builds in some common-sense typechecking that does this. Such issues epitomize the detail and precision that is required in computational modeling, and underscore the need to characterize coordination in a substantive, computational way.

4.4 Modeling collaborative agency

We spell out agents' deliberation in pursuit of successful collaborative conversation in terms of this background of information states and communicative intentions by specifying agents' cognitive architecture and the computational mechanisms of coordination that they use.

The architecture we use is shown in Figure 3. It instantiates a general cognitive architecture for perception, deliberation and action (Russell & Norvig 2003, Wooldridge 2000) but specializes it for collaborative agency by introducing symmetric processes of understanding and generation and synchronized processes of update. Such architectures are motivated in

```

loop {
  Perception
   $e \leftarrow \text{SENSE}()$ 
   $\langle \hat{c}, i \rangle \leftarrow \text{UNDERSTAND}(r, Z(c), e)$ 
   $c \leftarrow \text{UPDATE}(\hat{c}, i)$ 

  Determination and Deliberation
   $\hat{c} \leftarrow \text{ACT-TACITLY}(p, c)$ 
   $m \leftarrow \text{SELECT}(p, \hat{c})$ 
   $i \leftarrow \text{GENERATE}(r, \hat{c}, m, Z(c))$ 

  Action
   $\text{ACT-PUBLICLY}(a(i))$ 
   $c \leftarrow \text{UPDATE}(\hat{c}, i)$ 
}

```

Figure 3 A general, computational specification for collaborating agents that take both tacit actions and public actions.

greater detail in [Stone & Thomason 2003](#) and [Stone 2004b](#). The specification presented in [Figure 3](#) differs from earlier ones in allowing tacit actions, and maintaining the resulting uncertainties about context.

As always, it's best to explain the flow of control in this architecture by starting in the middle, with the deliberation that S does. When the deliberation begins, the state of the collaboration is represented as a checkpoint c . Now S has to decide how to proceed with the collaboration. To begin with, S determines and carries out whatever tacit actions are necessary at this stage of the collaboration, drawing on private information p .

The step:

$$\hat{c} \leftarrow \text{ACT-TACITLY}(p, c)$$

achieves this by updating the actual state of the collaboration from c to \hat{c} . Because S is the author these tacit actions, S knows that this update has taken place. However, these tacit actions introduce ambiguity for H and thus for the collaborators' mutual information. Before, the collaborators all mutually knew they were in state c ; now all they mutually know is that they are in some state from $Z(c)$.

Next S has to choose an appropriate public action to perform based on private information p and the actual state of the collaboration \hat{c} . This choice is accomplished by the step:

$$m \leftarrow \text{SELECT}(p, \hat{c}).$$

This choice results in a schematic “message” m that S intends to convey; the exact way that this action will be carried out must be chosen to preserve coordination at the next checkpoint, and m will therefore leave details open that will have to be specified later.

The final step of deliberation is generation. The generator takes a set of resources r which models the available communicative conventions in the domain. It also reasons from the message m and the actual context \hat{c} . It will need to construct a way of carrying out m that makes sense in \hat{c} . But the generator also has to flesh out this action enough to enable the hearer not only to recognize m but \hat{c} itself. That means the generator must take into account $Z(c)$, the horizon of alternatives that define the information state of the interaction. The generator now produces an intention i :

$$i \leftarrow \text{GENERATE}(r, \hat{c}, m, Z(c)).$$

This intention determines an action $a(i)$. Agent S now publicly carries out this action. Finally, to represent its effects, S updates the context from \hat{c} using intention i :

$$c \leftarrow \text{UPDATE}(\hat{c}, i).$$

This step implements agent S 's enlightened update, in relation to its own public action. Because S has taken a public action that makes S 's own tacit actions publicly recognizable, the new checkpoint context c incorporates the effects of both S 's recognizable tacit actions as well as the explicit contribution associated with the public action.

We now consider the beginning of the cycle of perception, deliberation and action from the standpoint of agent H . H is maintaining the initial checkpoint context as c , and so maintains the corresponding horizon $Z(c)$. H senses the action e that S has just made. H first understands e , using the communicative resources r and the contextual alternatives $Z(c)$:

$$\langle \hat{c}, i \rangle \leftarrow \text{UNDERSTAND}(r, Z(c), e)$$

The result has two parts. One is a representation of the context \hat{c} in which S must have been acting, reflecting the tacit actions that S implicitly carried out. The other is a representation i of the intention that motivated S 's action in the context. To reach the new checkpoint, H carries out the same update that S did:

$$c \leftarrow \text{UPDATE}(\hat{c}, i)$$

This context c is to be the next agreed checkpoint. Again, the UPDATE step has performed an enlightened update, capturing the effects both of the tacit actions that have occurred as well as the explicit contribution associated with the observed public action. It is now H 's turn to act.

This framework allows us to be precise about what is involved in maintaining coordination in dialogue. The requirements are formalized in (18), our PRINCIPLE OF COORDINATION MAINTENANCE.

- (18) For all resources r , checkpoints c , states $\hat{c} \in Z(c)$, and messages m :
 If i is obtained by $\text{GENERATE}(r, \hat{c}, m, Z(c))$,
 then $\langle \hat{c}, i \rangle = \text{UNDERSTAND}(r, Z(c), a(i))$.

The principle just says that, for any fixed perspective, any utterance $a(i)$ is understood as its speaker intends it. Suppose interlocutors share the background knowledge that fixes their communicative resources, their updates, and their horizons. Then if they share access to the initial conversational state, if they observe actions accurately, and if the coordination principle is satisfied, then they will continue to agree on each checkpoint.

This formal result points to the consistency of the kind of view we are advocating here. It shows that agents really could keep their view of context in alignment even without sharing all the information presupposed by their collaborators. Of course, the assumptions of the formal result are severe. Natural conversation succeeds even when there is more divergence in perspective and less certainty in action than our result considers. In this sense, the result calls attention to some of the idealizations and limitations in this provisional model. We leave more robust models for future treatments.

4.5 Coordination

The architecture of Figure 3 proposes substantive changes in how we think about and implement processes of understanding and generation to take tacit actions into account. Understanding no longer exploits a precisely

specified context; due to the horizon created by possible tacit actions, there is uncertainty. Conversely, generation needs to make a single context salient in a range of alternatives. Moreover, as we have seen, these new processes must coordinate their activities.

Despite these changes for understanding and generation, we can carry over a substantial amount of conceptual and computational infrastructure into the new framework. To pursue the parallel, let us make explicit the interpretive mechanisms one normally thinks about, where all the relevant information is known. Such mechanisms apply, in effect, to an individual candidate state c from the horizon. They deliver the set of maximally plausible interpretations for an action e in that state. We will use the notation $I(c, e)$ for this set of interpretations. The calculation of $I(c, e)$ is just interpretation with complete information. To carry out this calculation, we might use the abductive approach to interpretation advocated in Stone & Thomason 2003 and the incremental use of grammatical resources advocated in Stone 2004b. But there are certainly other ways to do it. (At the same time, we expect that an abductive approach to interpretation can also be extended to describe the understanding process in its entirety.)

As we hinted in Section 3, H must SIMULTANEOUSLY recognize which state from the horizon is actual and which interpretation action e has there. Here is a sketch of how this process might proceed.⁸ Agent H considers each state c in turn, computes $I(c, e)$ and inspects the result. It may be that $I(c, e)$ is empty. This means, for example, that action e would never be an appropriate continuation of the joint activity after the tacit actions that lead to c . Of course, this means that S could not have done those tacit actions; c can be eliminated. Alternatively, it may be that $I(c, e)$ has multiple elements. This means that e is ambiguous in c . This also means that c can be eliminated, because in c , S would not have fulfilled the obligations of coordination. If S had meant to convey one of these interpretations in c , he would have done e differently—for example, by adding words to disambiguate e . That leaves only states where $I(c, e)$ is unique. If there is exactly one such state c , then c must have been the actual state before the action was taken, and $I(c, e)$ must

⁸ In the probabilistic version of COREF (DeVault 2008, DeVault & Stone 2009), we have implemented a version of the context-sensitive understanding and generation processes sketched in this section. The implementation differs in a few ways from the discussion here, due to the presence of probabilistic information in the agent's reasoning about possible contexts and interpretations, and due to its use of simplified versions of the more principled reasoning outlined in this section; see DeVault 2008 for the details.

be the interpretation of the action there. Otherwise, since some sequences of tacit action are likely to be more obscure than others, it is convenient to assume a plausibility ordering over the states in the horizon. In this case, S must have meant the most plausible c for which $I(c, e)$ is unique. We'll use the notation $\text{BEST}(C)$ for the most plausible element of a set of states C .

We offer a formal summary of this discussion:

- (19) Suppose $\langle \hat{c}, i \rangle = \text{UNDERSTAND}(r, Z(c), e)$
 Then $\hat{c} = \text{BEST}\{q \in Z(c) : |I(q, e)| = 1\}$
 and $\{i\} = I(\hat{c}, e)$.

The formalism makes clear that we can, in principle, carry out interpretation in an uncertain context with the computational ingredients we use for interpretation in fully specified contexts.

Now let's consider how S plans an utterance. Suppose that S has a provisional utterance e in mind, and thereby intends to signal intention i with resources r in context \hat{c} and horizon $Z(c)$. Suppose, however, that $\langle \hat{c}, i \rangle$ is not what results from $\text{UNDERSTAND}(r, Z(c), e)$. Then e is not a satisfactory utterance and has to be modified somehow. There might be either of two problems with e . Perhaps $I(\hat{c}, e)$ is not uniquely $\{i\}$. This is the familiar kind of ambiguity that can come up even without tacit actions. The generator needs to signal interpretive connections that link \hat{c} to the intended interpretation i but that fail to link \hat{c} to other alternative interpretations that are at least as salient as i . The generator can do this by adding words to e and corresponding presupposed content to i . Alternatively, perhaps there is a more plausible state c' for which $I(c', e)$ is unique. This ambiguity arises only because of tacit actions and the possibility of enlightened update. Nevertheless, the solution for the generator is basically the same as always. The generator has to add words to e and corresponding presupposed content that links \hat{c} to the intended interpretation i but cannot be used to link c' to salient interpretations. In effect, just as we must distinguish i from its alternatives, we must distinguish \hat{c} from its alternatives too.

Again, the formalism makes clear that we can, in principle, carry out generation in an uncertain context with similar computational ingredients to those that fit fully specified contexts. As always we can build a generator by planning—constructing utterances and their intentions incrementally, using steps of grammatical derivation that add syntax, semantics and pragmatics simultaneously (Stone et al. 2003). By consulting models of interpretation and understanding, the generator can ensure that it conforms to the principles of

coordination required to align with others' understanding.

We can illustrate the kinds of reasoning we need in connection with dialogue (10). We have already seen that, in uttering *It's solid*, which advances the collaborative reference task by providing an additional constraint on the target object, COREF commits to the precondition that the previous clarification subdialogue is complete. This precondition is imposed by discourse coherence, not triggered by the grammar. However, COREF's commitment remains a presupposition on our account—and an informative one, since it might have been possible for COREF to continue the clarification subdialogue. COREF calculates that the hearer will in fact be able to recognize COREF's tacit completion of the clarification subdialogue from the kind of assertion that COREF is making. This is a complex prediction encapsulating a range of assumptions about the hearer. Had these assumptions been different, COREF would have chosen to say something more explicit here. For example, COREF could have used *The target square is solid*, which imposes additional grammatical preconditions and thereby distinguishes the target t_1 of the main collaborative reference task from the target S of the clarification subtask. In fact, we can see the effects of reasoning about coordination elsewhere in the utterances COREF does produce in (10). In generating C_3 , for example, COREF iteratively elaborates its description from *brown* to *light brown* so as to distinguish the light brown of its intended target from the darker brown also visible on the display.

5 Elaboration and discussion

We can now revisit our knowledge-level account of presupposition in light of the specific implementation we have provided for it. Here we draw on the ideas we have developed to offer some perspective on matters of ongoing debate in the literature on presupposition accommodation. Recall that we view presuppositions as ingredients of a speaker's communicative intention, and interpretation as interlocutors' coordinated recognition of this intention. In Section 5.1, we contrast this view with other models of how presuppositions are recognized, and suggest that our model clarifies the place of accommodation in interpretive reasoning. At the same time, in Section 5.2, we argue that it also clarifies the kinds of data that require substantive explanation as exemplifying exceptional patterns of interpretation. We suggest that our computational approach is a necessary ingredient in explaining how interlocutors manage to flout pragmatic rules as well as to follow them.

Of course, not all presuppositions involve planning, or depend on a well articulated context of conversational background and purposes. Our story remains provisional. We maintain only that if—as we believe—reasoning is a crucially important part of the story, then it is best to begin in a domain where reasoning is well behaved and can be formalized in detail. Then, if success is achieved in this domain, one attempts to extend the scope of the formalization. We suspect that, as in research on conversational implicature, too much attention has been paid to relatively exotic examples of presupposition accommodation, and that we might make better progress by concentrating on much larger bodies of everyday examples.

5.1 How accommodation takes place

David Lewis introduced a pretheoretic characterization of accommodation as an interpretive strategy for conversation.

- (20) If at time t something is said that requires presupposition P to be acceptable [or: something is said that requires for its truth that P be presupposed], and if P is not presupposed just before t , then—*ceteris paribus* and within certain limits—presupposition P comes into existence at t . (Lewis 1979: 340)

A wide range of research has now fleshed out this characterization by clarifying and formalizing aspects of the reasoning involved.

For example, it is important that this accommodation can be a natural and transparent adjustment to the state of the conversation. In effect, accommodation makes sure that an utterance is acceptable in the context in which it is interpreted. The process of understanding will show that P is required, and so will narrow down the alternatives to those that contain the information P . Thus when the hearer considers what response the utterance deserves, the utterance seems unproblematic. The fact that accommodation is in this sense an unreflective and inevitable side-effect of interpretation is emphasized by Stalnaker (1998) and von Stechow (2008).

Moreover, we can directly abstract a formal model of accommodation from Lewis's characterization. There are possibilities and limits for what can be accommodated, and this suggests that we should formalize the state of the conversation in terms of a SET OF ALTERNATIVES. We formalize the inference that P is required as narrowing down this set of alternatives to those that contain the information P . This is the influential and attractive picture of Beaver (2001).

These insights and formal tools largely carry over to our approach. For example, when an agent characterized as in Figure 3 plans or understands an utterance, it consults its horizon, a set of alternatives each of which paints a different picture of the state of the conversation. An utterance can select one of these states by imposing grammatical requirements that fit it alone. This mirrors Beaver's formalization. This reasoning recognizes the state of the conversation as a side-effect of the understanding process, much as Stalnaker and von Stechow describe it. In fact, as specified in Figure 3, the updates contributed by the utterance always advance to a state in which their preconditions are satisfied.

However, in our account, this reasoning is closely tied to an overall knowledge-level explanation of conversation as a collaborative interaction mediated by communicative intentions. The formal structures we use are independently motivated representations of communicative intention and the reasoning processes we use solve independently motivated problems of signaling and recognizing communicative intentions from utterances. Among other things, this means that when we apply our model to specify the behavior of a conversational agent or analyze examples of accommodation, we can work with a richer fabric of explanations. For instance, by thinking in terms of enlightened update, we can better understand how adjustments to the context as characterized by Lewis in Condition (20) can naturally arise for agents that coordinate successfully and collaborate effectively.

Take Stalnaker's and von Stechow's description of the evolution of presupposition as a side-effect of understanding. One motivation for this account is to save a theory of presupposition that takes a speaker's presuppositions to be part of the common ground. Since presuppositions will be part of the common ground by the time the utterance is updated, the argument goes, it's RIGHT to view the grammar as imposing requirements of mutuality.

We do not see how to incorporate this idea of sequential update into an account that does justice to general pragmatic reasoning. Since the interpreter must find the best overall explanation of the utterance, presupposition resolution can't be an isolated subsystem but must interact with other sources of alternative interpretations. Consider, for instance, the interaction of presupposition with ambiguity resolution. Suppose there are TWO competing interpretations, each of which imputes a different state to the conversation by the grammatical requirements it imposes. Each interpretation is consistent and unique in its own imputed context. The hearer cannot use the accommodated information to decide between them. Instead, the hearer

must fall back on the information that's ACTUALLY mutual—information from earlier in the conversation that creates a mutually evident preference for one context and its corresponding interpretation. The speaker, in formulating utterances, must do likewise. This is why we continue to offer a knowledge-level characterization of the problem of interpretation as drawing on mutual supposition.

Thus, while you can narrate our system's reasoning in Stalnaker's and von Stechow's terms, you CANNOT thereby eliminate the fact that there are TWO fundamentally irreconcilable bodies of information at play. There is the information that's ACTUALLY shared, as represented in our system by the horizon as a whole. And there is the information in the state of the conversation through which GRAMMATICAL REQUIREMENTS can be satisfied, as represented in our system by a single state on the horizon. Of course, that state on the horizon WILL be identified in utterance understanding, so all the information it contains WILL be mutually supposed before utterance update. But to develop a theory of presupposition and accommodation that explains interlocutors' successful coordination, both utterance planning and utterance understanding need to track the actual current body of mutual suppositions, rather than optimistically anticipating the results of accommodation.

Similarly, consider Beaver's model of accommodation as a process of eliminating alternatives for the context. Beaver motivates a set of alternative contexts by imagining discourse produced by an uncooperative speaker who is ignorant of the audience—so that the speaker's interpretation task resembles that of an overhearer, for whom the discourse is not intended. In such a circumstance, while the grammar may force the speaker to act as though certain information is mutually supposed, the audience will be uncertain about what information the speaker is appealing to. They can find this out from the grammatical requirements his utterance imposes.

Again, at the knowledge level, Beaver motivates a different reasoning problem from ours. Our system has no uncertainty about the common ground. It knows that what's mutual is ONLY what ALL states on the horizon have in common. It knows that information that's present in some states on the horizon but not in others is NOT mutually supposed. If imposing a grammatical requirement simply MEANS that information is mutually supposed, our system is WRONG to impose requirements that select some state on the horizon over others. It can't even be "flouting the rules," disobeying them to make a point; it doesn't have such a point. We characterize the problem our system solves differently. Our system, like any system that uses intentions

to pursue collaborations, represents intentions in part through commitments about the preconditions of intended actions. It makes those commitments when they are consistent and coordinates them based on mutual supposition.

5.2 The problematic of accommodation

All this being said, we acknowledge that accommodation in general, including cases that are harder to characterize as enlightened update, is in many ways a puzzling and challenging phenomenon. To deal with these cases, a theoretician may be driven to postulate pragmatic rules that are normally and routinely violated.

We can't help feeling that this is going too far—pursuing a methodology that makes the prospects for pragmatic theory seem rather dim. It is hard to make sense of rules that are normally and routinely violated. All the better, then, if we can account for a portion of the phenomena using abductive reasoning to the best explanation, where the explanation is a speaker plan. The larger the portion, the better for pragmatic theory. There may well be a residue, such as the case where Bess leaves the kitchen, Andrew breaks her crystal decanter, sweeps up the pieces, and announces when she returns:

(21) Don't worry, I'll replace your decanter with a better one.

But we can hope to keep this residue to a minimum.

In this section, we expand on these issues.

5.2.1 Accommodation on the cheap

Informative presupposition arises as a problem in the presence of a pragmatic rule requiring an utterance involving a presupposition to be appropriate only if its presuppositions are mutually supposed at that stage of the conversation. We are not committed to such a rule; the alternative rules that emerge from the ideas presented above would rather be (1) that an utterance involves a presupposition *P* if the intention underlying the utterance requires a commitment to a precondition *P* and (2) that an utterance is only appropriate to the extent that its presuppositions can be recognized and added to the common ground.

To take something on board in a conversation is not the same as to believe it. As [Stalnaker \(1975\)](#) points out, the appropriate modality is “supposing for

the sake of the conversation.” A hearer can be credulous, coming away from a conversation fully believing it all, or can be skeptical and selective.

Often, we may be willing to take something on board in this way to avoid distracting the flow of the conversation. Consider the example of the ailing cat once again. You call me and say:

(22) I need to cancel our lunch meeting. I have to take my cat to the veterinarian.

On hearing this, I may be mildly surprised. In fact, I thought you were a dog lover. But the point of the conversation is our lunch meeting, and—even if I suspect you of making up an excuse—I would rather not confront you or call for a detailed explanation. I get on with the task of rescheduling the lunch, thereby taking the hypothetical cat on board for the sake of the conversation.

Even if I’m not at all surprised or suspicious—you are, let’s suppose, the sort of person I’d expect to own a cat—what I take on board at this point is (1) irrelevant to the main topic, and (2) underspecified. In effect, all I am adding are two almost naked discourse referents: a cat belonging to you and a veterinarian who is attending to the cat. They are not relevant to the topic, are not likely to be mentioned again, and I can forget them more or less immediately if I care to. The effort involved in this accommodation is minimal.

On the other hand, when information is topical and important, and certainly when it is controversial, it is harder to slip it in by accommodation. If you call me asking me to pick up your cat at the veterinarian, you are unlikely to begin the conversation by saying

(23) Hello. Can you pick up my cat at the vet?

if I have no idea you have a cat. It is politer and more orderly to explicitly introduce the cat and its vicissitudes before making the request.

For accommodation to work well, it should be cognitively inexpensive.

5.2.2 The meaning of presupposition failure

But there are exceptions. For example, it is fairly common for the first sentence of a book to have presuppositions, and to violate normal rules about the orderly presentation of information. The first sentence of *The*

Missing World,⁹ for instance, is this:

- (24) They were quarrelling on the phone when it happened, although anyone overhearing them might easily have failed to detect the fury that lay behind their pragmatic sentences.

In the case of a novel, cases like this do not seem to be cognitively expensive, although they can be annoying if the author doesn't follow through by explaining the context. Nor do they seem to exemplify the kind of enlightened update we have explored, where speakers appeal to readily identified real-world information to contribute naturally to discourse. Instead, examples of this kind are motivated and shaped by stylistic considerations, and intended to achieve an artistic effect. And they are genre specific; you would not expect a conversation or a tractor manual to begin with a sentence that takes the liberties that Sentence (24) takes.

These considerations suggest that we must distinguish between enlightened update and other cases where a speaker uses an utterance whose presupposition *P* is not mutually known. One possibility is that the speaker pretends to be somebody else, or to be addressing some other audience than is actually present, as in Clark's (1992) pragmatic reconstruction of storytelling, humor or irony in terms of "layering." In effect, then, the presupposition *P* characterizes the state of the pretend discourse, and the interpreter must recognize the pretense in order to understand the discourse. That seems to fit Example (24) from Livesey's novel.

Such explanations may require us to change the way we think of the problem of interpretation as a whole, not just the process of presupposition accommodation. Now interpretation involves the open-ended reconstruction of the speaker's pretense, and allows the audience to explore a wide range of assumptions to make sense of the utterance. Even hypothetical alternatives to the preceding discourse can be considered. In fact, this is a common feature of jokes (Richie 2003). And it may explain the pronouns in Example (24). Moreover, what we really take on board from a joke or story is only indirectly related to the context we reconstruct for it. Perhaps the assumptions give voice to a real truth that we would not otherwise broach. Or perhaps the assumptions provide a foil to the way we think the world is. So the reasoning we need here is not just a matter of accommodating the speaker's presupposition. In many of these cases, only knowledge about the speaker can settle which, if any, was intended. In others, the intention

⁹ By Margot Livesey. Alfred A. Knopf, New York, 2000.

may be incomplete, artificial, and obscure, and the problem of interpretation will become open-ended, conscious, and problematic. This sort of difference between literary art and tractor manuals is to be expected.

Karttunen's historically important example (a commencement announcement) is a similarly outlying case, in which the presupposition is the new information that the author intends to communicate.

- (25) We regret that children cannot accompany their parents to commencement exercises.

Cases like these are more like appositive constructions than like ordinary presupposition sentences—that is, sentence (25) is like

- (26) Children, we regret, cannot accompany their parents to the commencement exercises.

We suspect that such cases represent constrained and perhaps even conventionalized discourse strategies. It may be that these usages can be explained by a desire to save face or evade responsibility. By using an utterance which (exceptionally) presupposes *P* to say that *P*, the speaker avoids the need to overtly to raise *P* as a questionable topic of discussion (Thomason 1990). Moreover, the kinds of presuppositions that can be made mutually available through this strategy, while perhaps broader than those that result from tacit actions, remain quite limited; for example, one cannot retrospectively change what utterances have been made in the conversation (Beaver & Zeevat 2007). Corpus evidence would be one way to confirm our guess that usages like these are stylistically marked, are confined to a few “presupposition triggers,” and tend to be used when the speaker wants to evade an issue.

5.2.3 Two knowledge levels

How then should we think of exceptions? Perhaps we need to distinguish between modeling the reasoning itself and giving a rational reconstruction of the principles that make the reasoning coherent.

Allan Newell's “The knowledge level” was the first AAAI Presidential address, delivered at Stanford University in 1980.

Newell took as his starting point the emergence within computer science of a number of levels of from which it could be fruitful to describe a computational system, with each level characterized by its own ontology and laws. He lists five levels, including the CIRCUIT LEVEL and the SYMBOLIC or PROGRAM LEVEL. While a computer engineer might use the former level (ignoring

the program entirely and concentrating on the flow of digital information through the registers of the device) to analyze the behavior of a computer running a simple program, a software engineer, taking the correct behavior of the hardware for granted, might concentrate on the transitions between symbolic states that are induced by the program.

To these levels, Newell suggested the addition of a higher level, involving the attribution of BELIEFS and GOALS to the system. This idea is congenial to computer scientists, since it is natural and unavoidable to think even of simple systems as having goals and beliefs. For instance, it is natural to explain why my navigation system is beeping by saying it thinks I have reached a turn in my route and wants to alert me to turn. But the idea has also been adopted by some philosophers; see, for instance, [Dennett 1979](#).

In his later work, especially in [Newell 1992](#), Newell elaborated these ideas and argued for their general applicability in cognitive science. There, he also speculates about other levels, including a social level. He doesn't provide details about the sort of laws that would be appropriate for the social level, but presumably (like the knowledge level, which could make use of a variety of formalization devices, including those of probability and logic) the social level could be formalized using ideas from macroeconomics, or using a logic providing for the formalization of group norms.

This distinction can be useful in thinking about accommodation. Anyone who takes accommodation or similar phenomena seriously is likely to be plagued by an apparent appeal to a rule that is then undermined by the accommodating practice. (And it is hard not to take it seriously as a general phenomenon, on seeing the range of examples that Lewis lays out in [Lewis 1979](#).) The undermining, which affects Lewis' account in [Lewis 1979](#), as well as Grice's idea, formulated in [Grice 1989](#): 30ff, of how implicatures are induced by "flouting" a rule of conversation, conforms to the following generic pattern.¹⁰

- (27) A general rule R of the form *Do A only if C* is postulated, where C is a condition that can be manipulated by an audience H . An agent S is observed by H to be doing A while C is mutually known to be false. H then acts to fulfil C , and S expects H to so act.

To take a nonlinguistic example, a woman at a restaurant violates the rule (R₁) "Don't sit down unless there is an empty chair positioned behind

¹⁰ This formulation is meant to include cases where the audience acts to make a condition true, and cases where the audience simply supposes the condition for the sake of the conversation.

you,” knowing that a waiter will position the required chair, and that indeed her sitting will be the signal for this act. Here, it seems natural to say either (taking a nonmonotonic view of the rule) that R1 has not been violated, because it is defeasible and this is one of the exceptional conditions, or (taking a monotonic view) that R1 has to be replaced with an improved, qualified version: “Don’t sit down unless there will be an empty chair positioned behind you when it is needed.”

Von Fintel uses the monotonic tactic in [von Fintel 2008](#) to defend Stalnaker’s rule (R2) “Do not trigger a presupposition *P* unless *P* is in the common ground that is available for updating the conversation.” The idea is that, like the waiter, the audience can add the required presupposition to the common ground before updating with the content of the utterance.

Von Fintel seems to treat pragmatic rules like Stalnaker’s as indefeasible—that is, he replaces R2 with a more accurate, qualified rule. But a defeasible approach to pragmatic rules might be more explanatory. In the restaurant scenario, it is natural to say that the woman is acting as if there were a chair behind her. But it is hard to explain how she could be doing this unless R1 has some force as a rule of conduct. Also, giving R1 some force as a *ceteris paribus* rule can also provide a more plausible reconstruction of the woman’s and the waiter’s reasoning. If someone were to feel faint in a restaurant, and involuntarily begin to collapse into a sitting position, a waiter might well rush to the spot with a chair. But in the accommodation scenario, we have a strong intuition that the woman MEANS SOMETHING by sitting, and that the waiter has correctly recognized this meaning. If we want to provide a rational reconstruction of this reasoning (rather than treating it as a special convention that applies in some restaurants), it will be useful to give rulelike force to R1.

Similarly, taking R2 to be defeasible could help to explain Stalnaker’s claim that when *S* says “I have to pick up my cat at the veterinarian,” *S* is acting as if it were part of the common ground that *S* has a cat.¹¹

But the analogy between accommodation in conversation and in the restaurant, and attempts like von Fintel’s to neutralize the undermining effect of accommodation, is imperfect in cases in which accommodation leads to routine violations of the rule. As many people have pointed out,

¹¹ In [Stalnaker 1975: 202](#), Stalnaker suggests that a speaker is pretending that the presupposition is part of the common ground in such cases. Although there are cases, like (24), that are best described as pretenses, Stalnaker’s cat example (3) is not one of these. Here, it is probably better to say that the speaker is acting as if it is part of the common ground.

violations of R2 occur with great frequency, and there is no intuition in such cases that any rule or expectation has been violated. This makes it difficult to give R2 the status of a defeasible rule—at least, if we expect defeasible rules to hold usually, or normally. The rule we need to invoke to explain the accommodation really seems to be undermined when accommodation becomes the rule.

Here it may be useful to think of R2 not as a knowledge-level rule governing behavior, but as a rule at the social level. It is possible for a rule to have acknowledged status as a social norm of a fairly idealized sort, and yet for it to be routinely flouted without any sense of violation. “Don’t interrupt” might be an example, in contemporary American society. If we assume a society of accommodating agents, who are also capable of learning and forming habits, we would expect the rules that give rise to instances of accommodation, over time, to give rise to habits that flout these rules. But this doesn’t mean that these rules will not preserve a vestigial status, or that they won’t be useful for certain theoretical purposes.

Social rules can certainly be violated in order to achieve an effect—think of refusing to shake a hand that is offered. But when the effect is cooperative and matches a familiar and habitual pattern, the perceived effect is quite different, and it is harder to imagine that a rule has been violated.

We remain neutral on whether such rules can play a useful and legitimate role in pragmatic theory. But invocations of idealized social rules that are routinely flouted have figured in pragmatic rational reconstructions, such as Grice’s rationale for implicatures and the reconstruction that [Stalnaker \(1998\)](#) offers of Lewis’s account of presupposition accommodation. Perhaps vestigial or idealized rules could play a useful explanatory role in a systematic theory of pragmatic norms—providing coherence, motivating or even enabling the derivation of lower-level rules, and explaining how these rules are learned. However, since we lack a systematic theory of this kind, these suggestions are only tentative.

Another alternative is to look for improved pragmatic rules, even at the social level. If, as we suggested, the common ground is a complex intentional structure and the propositions that figure in it may play different roles, then treating the common ground itself as a simple set of propositions may be an oversimplification—there may be more than one way to add a proposition to the common ground. Independently, and for quite different reasons, some pragmatic theorists have made similar suggestions; see, for instance, [Horn 2007](#).

5.2.4 Explanation and implementation

In any case, it would be a mistake to invoke rules that are “made to be broken” in a knowledge-level specification of problem solving in conversation. A conversational planner S has the goal, say, of updating the common ground with an individual c , which is a cat and belongs to S , with an individual v who is a veterinarian, and with the proposition that S has to pick up c at v 's office. There are various ways of accomplishing this goal, including:

- (28) I have a cat.
The cat is at a veterinarian.
The veterinarian has an office.
I have to pick up the cat at the veterinarian's office.
- (29) I have a cat.
I have to pick up the cat at the veterinarian's.
- (30) I have to pick up my cat at the veterinarian's.

The planner needs to know that such alternatives will accomplish the same result, and to know that normally the last version is the preferred, unmarked way of doing this. But to reconstruct the planner's knowledge and reasoning, we do not need to postulate a rule prohibiting a definite reference to an individual not already in the common ground. Indeed, such a rule could quite likely get in the way of using the unmarked formulation. A rational reconstruction is not the same thing as a knowledge-level specification.

In the recent literature on presuppositions, it has been suggested that many presuppositions arise as conversational implicatures; see, for instance, [Simons 2001](#). We do not need to take a stand on this issue here, as long as Simons's suggestion is thought of as a high-level explanation of presupposition, perhaps at the social level. At the implementation level, and even perhaps at the knowledge level, it would be difficult or impossible to account for presuppositions by reasoning processes that appealed to first principles rather than specific grammatical conventions and specific contextual resources. Implicatures can in principle be derived from first principles using abductive reasoning, but the knowledge that would be needed to support the processes is very open-ended. A general solution to the problem of conversational implicature would probably require robust, general-purpose common-sense reasoning. In any case, even if many or even all presuppositions originate as conversational implicatures, many of them would be highly *generalized* implicatures, ones that hold over a large set of interpretation contexts. In

many reasoning domains, it would be a reasonable idealization to treat these implicatures as having a grammatical origin.

6 Concluding Remarks

In this paper, we have developed an account of presupposition and accommodation grounded in a computational model of reasoning in conversation. We have focused on cases of ENLIGHTENED UPDATE, in which the conversational state advances to reflect the speaker's recognizable private commitments as well as the explicit contribution of an utterance. Our formal model shows how enlightened update arises naturally from the reasoning that interlocutors need in conversation and other collaborative activities. It also shows how enlightened update offers a powerful explanation for many otherwise puzzling cases of informative presupposition (though not for all of them).

Our work ties together linguistic, formal and philosophical approaches to presupposition and accommodation. Much work remains to reconcile the insights of these different traditions. We close with an overview of three challenges that we think are particularly pressing (drawing in part on critical discussion from the workshop on presupposition accommodation at the Ohio State University in 2006 where we presented an earlier version of this paper). We see these issues as setting an agenda for further research in computational approaches to semantics and pragmatics.

6.1 Strategic reasoning and varieties of interpretation

Methodologically, we are committed to precise and implementable accounts of pragmatic reasoning. Our model of enlightened update characterizes a specific kind of interpretive reasoning. Our discussion in Section 5.2 speculates that other kinds of reasoning will be required for some of the cases of informative presupposition discussed in the literature, because different kinds of strategic considerations inform language users' interpretations.

This contrasts with a common situation in pragmatics, where researchers propose to unify disparate interpretive effects through broad, programmatic principles. Grice (1989) has the Cooperative Principle, Sperber & Wilson (1995) have the Principle of Relevance, and even Lewis's (1979) rule of accommodation is quite general in its application. We have argued that the uniformity these approaches profess could not be sustained under more precise and systematic formalization.

To defend this claim, however, we depend on the results of ongoing empirical work in linguistics that promises to give a better handle on the diversity of interpretation. At the same time, we depend on computational work that can scale up abductive models of interpretation. Abductive models predict that different kinds of content require different interpretive assumptions which impact overall preferences for interpretation. Any assumptions must also be reconciled with an interlocutor's understanding of the conversational situation and of the speaker. So connecting our theory with a broader range of linguistic evidence fundamentally requires systematic ways of articulating the expectations and preferences that interlocutors bring to richer conversations. This remains a challenge for AI research.

6.2 Communicative intentions and compositional semantics

The point of planning in AI is to allow agents to formulate and carry out complex sequences of actions in order to achieve a goal. When a plan specifies that an agent should carry out first one action and then a second, the plan must guarantee that the preconditions of the second action will be satisfied in the situation that results from the execution of the first. Otherwise the second action will fail. An intention to carry out a complex plan therefore incorporates appropriate conditional commitments about how the world will be in the future if certain actions are taken.

These observations show how our ideas connect to PRESUPPOSITION PROJECTION, the problem of predicting the presuppositions inferred from uses of grammatically complex utterances. We can think of such utterances as the outgrowth of incremental plans to contribute information to the conversation constituent-by-constituent or even word-by-word. If any of the component actions have preconditions, then the plan must show how those preconditions will be satisfied at the point the action is taken. In general, then, presuppositions will describe not only the initial context in which the overall utterance is made but also the local contexts established incrementally as each constituent effects its own update to the conversation in turn. This aligns with the familiar principle that the presuppositions triggered by an item must be satisfied in that item's local context (Heim 1983).

Unfortunately, the details are problematic. Planning theory demands a concrete model of utterances as actions, in which the atomic units of semantic composition transparently trigger atomic changes in conversational state. But traditional dynamic semantics does not restrict itself to such

operations (Groenendijk & Stokhof 1991). For example, it handles universal quantifications by performing the update associated with the nuclear scope repeatedly, once for each object in the restrictor set. No such repeated updates actually occur as the utterance is performed—the nuclear scope arrives just once in the performance of the utterance. A planning theory must likewise complete the meaning of the universal through a single update.

In fact, researchers are increasingly developing theories of local context such as Schlenker's (2010) which are grounded in incremental specification of utterances rather than in merely metaphorical dynamics. We are optimistic that such approaches will substantiate our analysis of speaker presupposition in terms of communicative intentions and enable it to be extended naturally to the compositional specification of complex utterances.

6.3 Building agents with communicative intentions

Philosophical intuitions single out communicative intentions in part by their content and in part by the role that they play in agents' cognitive economy. Our work in COREF focuses on two features of communicative intentions: they must represent preconditions, and they must play their role in agents' coordination with one another. There is much more to be said about intentions in general, and communicative intentions in particular. COREF is not yet responsive to a range of desiderata that we would ultimately expect in a formal model of collaborative rationality in conversation.

To start, there are problems of content. Communicative intentions need to track the specific mechanisms that distinguish communication from other ways of getting information across (Grice 1957). Moreover, they must show how agents defer to usage in the broader community and fit their meanings together across a conversation (Putnam 1975). Such details are crucial to our intuitions about meaning, but COREF's intentions do not yet represent them. Even the representations COREF does have are preliminary. It is not clear, for example, how to map out a consistent and intuitively defensible set of assumptions about the state of a conversation where one interlocutor has failed to understand something the other has said, as in the clarification subdialogue considered in this paper.

More difficult questions surround the role intentions play in COREF's deliberation. Intention, in Cohen & Levesque's (1990) insightful slogan, is choice plus commitment. COREF does not currently act as though it is fully committed to its communicative intentions, in all the ways our intuitions

demand. For example, while COREF can use its intentions to follow up or clarify what it has said, it cannot use them as people do to identify differences in perspective, resolve possible misunderstandings and negotiate a shared understanding of the distinctions at issue in a conversation. Without these abilities, it is difficult to understand COREF's representations as communicative intentions (DeVault, Oved & Stone 2006).

Intentions generally play many other roles. They can focus efforts in planning for the future by ruling out options that are incompatible with decisions already made (Bratman 1987). Attribution of intentions may also play a role in explaining one's own or others' actions retrospectively, to interpret or critique problematic patterns of behavior (Suchman 1987). COREF's communicative intentions do not play these roles—but we are not convinced that they should. Future models of collaborative agency need to pay clearer attention to the philosophical distinction between intentions-in-action, such as the communicative intention behind a specific utterance, and future-oriented intentions, of the sort that guide extended patterns of collaboration and inquiry (Searle 1983). Of particular importance is the need to characterize the functional interplay between the two kinds of attitudes.

Despite the many challenges that remain in building agents with communicative intentions, we are optimistic that the key principles we have presented will be lasting ingredients of the designs. For example, it is possible to do better justice to the semantic relationships that mediate extended patterns of inquiry in conversation by using presuppositions of coreference to ensure that related utterances address the same issues. These expressive representations can support more flexible deliberation within the same kinds of computational architecture we have considered here. DeVault et al. (2006) suggest that realizing such deliberation might enable response to criticisms of COREF's its communicative intentions.

Of course, actually building such systems inevitably leads to unexpected issues and refinements to the theory. The fact that such ideas can be tested in this way is, we believe, an important contribution of AI methodology to pragmatics, an area which desperately needs more effective ways of confirming and disconfirming theories.

References

- Abusch, Dorit. 2010. Presupposition triggering from alternatives. *Journal of Semantics* 27(1). 37–80. doi:10.1093/jos/ffp009.
- Aumann, Robert J. 1976. Agreeing to disagree. *Annals of Statistics* 4(6). 1236–1239. doi:10.1214/aos/1176343654.
- Barker, Chris. 2002. The dynamics of vagueness. *Linguistics and Philosophy* 25(1). 1–36. doi:10.1023/A:1014346114955.
- Beaver, David. 2001. *Presupposition and assertion in dynamic semantics*. Stanford, California: CSLI Publications.
- Beaver, David I. & Henk Zeevat. 2007. Accommodation. In Gillian Ramchand & Charles Reiss (eds.), *Oxford handbook of linguistic interfaces*, 503–538. Oxford: Oxford University Press.
- Bratman, Michael E. 1987. *Intentions, plans and practical reason*. Harvard University Press.
- Brennan, Susan E. & Herbert H. Clark. 1996. Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory and Cognition* 22(6). 1482–1493. doi:10.1037/0278-7393.22.6.1482.
- Burge, Tyler. 1973. Reference and proper names. *The Journal of Philosophy* 70(14). 425–439. doi:10.2307/2025107.
- Clark, Herbert. 1992. *Arenas of language use*. Chicago: University of Chicago Press.
- Clark, Herbert H. & Catherine R. Marshall. 1981. Definite reference and mutual knowledge. In Arivind Joshi, Bonnie Webber & Ivan Sag (eds.), *Elements of discourse understanding*, 10–63. Cambridge, England: Cambridge University Press.
- Clark, Herbert H. & Deanna Wilkes-Gibbs. 1986. Referring as a collaborative process. *Cognition* 22(1). 1–39. doi:10.2307/2025107.
- Cohen, Philip R. & Hector J. Levesque. 1990. Intention is choice with commitment. *Artificial Intelligence* 42(2–3). 213–261. doi:10.1016/0004-3702(90)90055-5.
- Dale, Robert & Ehud Reiter. 1995. Computational interpretations of the Gricean maxims in the generation of referring expressions. *Cognitive Science* 19(2). 233–263. doi:10.1207/s15516709cog1902_3.
- Dennett, Daniel C. 1979. Artificial intelligence as philosophy and as psychology. In Martin Ringle (ed.), *Philosophical perspectives in artificial intelligence*, 57–80. Atlantic Highlands, New Jersey: Humanities Press.
- DeVault, David. 2008. *Contribution tracking: Participating in task-oriented*

- dialogue under uncertainty*. New Brunswick, New Jersey: Department of Computer Science, Rutgers, The State University of New Jersey dissertation.
- DeVault, David, Natalia Kariaeva, Anubha Kothari, Iris Oved & Matthew Stone. 2005. An information-state approach to collaborative reference. In *ACL 2005 proceedings companion volume: Interactive poster and demo track*, 1-4.
- DeVault, David, Iris Oved & Matthew Stone. 2006. Societal grounding is essential to meaningful language use. In *Twenty-first national conference on artificial intelligence (AAAI-06)*, 747-754.
- DeVault, David, Charles Rich & Candace L. Sidner. 2004. Natural language generation and discourse context: Computing distractor sets from the focus stack. In Valerie Barr & Zdravko Markov (eds.), *17th international FLAIRS conference (FLAIRS-2004)*, Menlo Park, California: AAAI Press.
- DeVault, David & Matthew Stone. 2006. Scorekeeping in an uncertain language game. In *Proceedings of the 10th workshop on the semantics and pragmatics of dialogue (SemDial-10)*, 139-146.
- DeVault, David & Matthew Stone. 2007. Managing ambiguities across utterances in dialogue. In *Proceedings of the 11th workshop on the semantics and pragmatics of dialogue (Decalog 2007)*, 49-56.
- DeVault, David & Matthew Stone. 2009. Learning to interpret utterances using dialogue history. In *Proceedings of the 12th conference of the European Association for Computational Linguistics (EACL)*, .
- Fagin, Ronald, Joseph Y. Halpern, Yoram Moses & Moshe Y. Vardi. 1995. *Reasoning about knowledge*. Cambridge, Massachusetts: The MIT Press.
- von Fintel, Kai. 2008. What is presupposition accommodation, again? *Philosophical Perspectives* 22(1). 137-170. doi:10.1111/j.1520-8583.2008.00144.x.
- Gardent, Claire & Laura Kallmeyer. 2003. Semantic construction in feature-based TAG. In *10th meeting of the European chapter of the Association for Computational linguistics*, 123-130.
- Geanakopulos, John. 1994. Common knowledge. In Robert Aumann & Serigiu Hart (eds.), *Handbook of game theory, with economic applications, vol. 2*, chap. 40. Amsterdam: North-Holland.
- Grice, H. Paul. 1957. Meaning. *Philosophical Review* 66. 377-388. doi:10.2307/2182440. Republished in H.P. Grice, *Studies in the Way of Words*, Harvard University Press, 1989.
- Grice, H. Paul. 1989. *Studies in the way of words*. Cambridge, Massachusetts: Harvard University Press.

- Groenendijk, Jeroen & Martin Stokhof. 1991. Dynamic predicate logic. *Linguistics and Philosophy* 14(1). 39-100. doi:10.1007/BF00628304.
- Gundel, Jeanette K., Nancy Hedberg & Ron Zacharski. 1993. Cognitive status and the form of referring expressions in discourse. *Language* 69(2). 274-307. <http://www.jstor.org/stable/416535>.
- Heeman, Peter & Graeme Hirst. 1995. Collaborating on referring expressions. *Computational Linguistics* 21(3). 351-382. <http://aclweb.org/anthology-new/J/J95/J95-3003.pdf>.
- Heim, Irene. 1983. On the projection problem for presuppositions. In *Second annual West Coast conference on formal linguistics*, 114-126.
- Hobbs, Jerry, Mark Stickel, Douglas Appelt & Paul Martin. 1993. Interpretation as abduction. *Artificial Intelligence* 63(1-2). 69-142. doi:10.1016/0004-3702(93)90015-4.
- Horn, Laurence R. 2007. Toward a Fregean pragmatics: *Voraussetzung, Nebengedanke, Andeutung*. In Istvan Kecskes & Laurence R. Horn (eds.), *Explorations in pragmatics: Linguistic, cognitive and intercultural aspects*, 39-69. Berlin: Mouton de Gruyter.
- Joshi, Aravind K., Leon S. Levy & Masako Takahashi. 1975. Tree adjunct grammars. *Journal of the Computer and System Sciences* 10(1). 136-163. doi:10.1016/S0022-0000(75)80019-5.
- Joshi, Arivind. 1982. Mutual beliefs in question-answer systems. In N.V. Smith (ed.), *Mutual knowledge*, 181-197. London: Academic Press.
- Kaplan, David. 1989. Demonstratives: an essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In Joseph Almog, John Perry & Howard Wettstein (eds.), *Themes from Kaplan*, 481-563. Oxford: Oxford University Press.
- Kripke, Saul. 1972. Naming and necessity. In Gilbert Harman & Donald Davidson (eds.), *Semantics of natural language*, 253-355. Dordrecht: D. Reidel Publishing Co.
- Larsson, Staffan & David Traum. 2000. Information state and dialogue management in the TRINDI Dialogue Move Engine Toolkit. *Natural Language Engineering* 6. 323-340. doi:10.1017/S1351324900002539.
- Lewis, David K. 1969. *Convention: A philosophical study*. Cambridge, Massachusetts: Harvard University Press.
- Lewis, David K. 1979. Scorekeeping in a language game. *Journal of Philosophical Logic* 8(3). 339-359. doi:10.1007/BF00258436.
- Lochbaum, Karen E. 1998. A collaborative planning model of intentional structure. *Computational Linguistics* 24(4). 525-572. <http://aclweb.org/>

- [anthology-new/J/J98/J98-4001.pdf](#).
- Mackworth, Alan. 1987. Constraint satisfaction. In S.C. Shapiro (ed.), *Encyclopedia of artificial intelligence*, 205-211. John Wiley and Sons.
- Marr, David. 1982. *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: W.H. Freeman.
- Newell, Allen. 1981. The knowledge level. *Artificial Intelligence Magazine* 2(2). 1-20. <http://www.lirmm.fr/~cerri/teachingo2/newell/knowledge.pdf>, <http://www.aaai.org/Resources/Papers/AIMag13-04-002.pdf>.
- Newell, Allen. 1992. *Unified theories of cognition*. Cambridge, Massachusetts: Harvard University Press.
- Perrault, C. Raymond & James F. Allen. 1980. A plan-based analysis of indirect speech acts. *American Journal of Computational Linguistics* 6(3-4). 167—182. <http://aclweb.org/anthology-new/J/J80/J80-3003.pdf>.
- Poesio, Massimo & David R. Traum. 1997. Conversational actions and discourse situations. *Computational Intelligence* 13(3). 309-347. doi:10.1111/0824-7935.00042.
- Pollack, Martha. 1990. Plans as complex mental attitudes. In Philip R. Cohen, Jerry Morgan & Martha Pollack (eds.), *Intentions in communication*, 77-103. Cambridge, Massachusetts: MIT Press.
- Pollack, Martha. 1992. The uses of plans. *Artificial Intelligence* 57(1). 43-68. doi:10.1016/0004-3702(92)90104-6.
- Purver, Matthew. 2004. *The theory and use of clarification requests in dialogue*. London: King's College, University of London Ph.d. dissertation. Available at <http://www.stanford.edu/~mpurver/papers/purvero4thesis.pdf>.
- Putnam, Hilary. 1975. *Mind, language and reality: Philosophical papers, vol. 2*. Cambridge, England: Cambridge University Press.
- Reiter, Raymond. 2001. *Knowledge in action: Logical foundations for specifying and implementing dynamical systems*. Cambridge, Massachusetts: The MIT Press.
- Rich, Charles, Candace L. Sidner & Neal Lesh. 2001. COLLAGEN: Applying collaborative discourse theory to human-computer interaction. *Artificial Intelligence Magazine* 22(4). 15-26.
- Richie, Graeme. 2003. *The linguistic analysis of jokes*. London: Routledge.
- Roberts, Craige. 2003. Uniqueness in definite noun phrases. *Linguistics and Philosophy* 26(3). 287-350.
- Russell, Stuart & Peter Norvig. 2003. *Artificial intelligence: A modern approach*. Englewood Cliffs, New Jersey: Prentice Hall 2nd edn.

- Schabes, Yves. 1990. *Mathematical and computational aspects of lexicalized grammars*: Computer Science Department, University of Pennsylvania dissertation.
- Schlenker, Philippe. 2010. Presuppositions and local contexts. *Mind* 119. 377-391. doi:10.1093/mind/fzq032.
- Searle, John R. 1983. *Intentionality: An essay in the philosophy of mind*. New York: Cambridge University Press.
- Simons, Mandy. 2001. On the conversational basis of some presuppositions. In *Semantics and linguistic theory XI*, 431-448. Available at <http://www.hss.cmu.edu/philosophy/simons/simons.SALT2001.pdf>.
- Simons, Mandy. 2003. Presupposition and accommodation: Understanding the Stalnakerian picture. *Philosophical Studies* 112(3). 251-278. doi:10.1023/A:1023004203043.
- Sperber, Dan & Deirdre Wilson. 1995. *Relevance, second edition*. Oxford: Blackwell.
- Stalnaker, Robert. 1973. Presuppositions. *Journal of Philosophical Logic* 2(4). 447-457. doi:10.1007/BF00262951.
- Stalnaker, Robert. 1975. Pragmatic presuppositions. In Milton K. Munitz & Peter Unger (eds.), *Semantics and philosophy*, 197-213. New York: Academic Press.
- Stalnaker, Robert. 1978. Assertion. In Peter Cole (ed.), *Pragmatics: Syntax and semantics* 9, 315-332. New York: Academic Press.
- Stalnaker, Robert. 1998. On the representation of context. *Journal of Logic, Language, and Information* 7(1). 3-19. doi:10.1023/A:1008254815298.
- Stalnaker, Robert. 2002. Common ground. *Linguistics and Philosophy* 25(5-6). 701-721. doi:10.1023/A:1020867916902.
- Stickel, Mark E. 1991. A Prolog-like inference system for computing minimum-cost abductive explanations in natural-language interpretation. *Annals of Mathematics and Artificial Intelligence* 4. 89-105. doi:10.1007/BF01531174.
- Stone, Matthew. 2004a. Communicative intentions and conversational processes in human-human and human-computer dialogue. In John Trueswell & Michael Tanenhaus (eds.), *World situated language use: Psycholinguistic, linguistic, and computational perspectives on bridging the product and action traditions*, Cambridge, Massachusetts: The MIT Press.
- Stone, Matthew. 2004b. Intention, interpretation and the computational structure of language. *Cognitive Science* 28(5). 781-809. doi:10.1207/s15516709cog2805_7.
- Stone, Matthew, Christine Doran, Bonnie Webber, Tonia Bleam & Martha

- Palmer. 2003. Microplanning with communicative intentions: The SPUD system. *Computational Intelligence* 19(4). 311–381. doi:10.1046/j.0824-7935.2003.00221.x.
- Stone, Matthew & Richmond H. Thomason. 2002. Context in abductive interpretation. In Johan Bos, Mary Ellen Foster & Colin Matheson (eds.), *EDILOG 2002: Proceedings of the sixth workshop on the semantics and pragmatics of dialogue*, 169–176. Edinburgh: Cognitive Science Centre, University of Edinburgh.
- Stone, Matthew & Richmond H. Thomason. 2003. Coordinating understanding and generation in an abductive approach to interpretation. In Ivana Kruijff-Korbayová & Claudia Kosny (eds.), *Diabruck 2003: Proceedings of the 7th workshop on the semantics and pragmatics of dialogue*, 131–138. Saarbrücken: Universität des Saarlandes.
- Suchman, Lucy A. 1987. *Plans and situated actions: the problem of human-machine communication*. New York, NY, USA: Cambridge University Press.
- Thomason, Richmond H. 1990. Accommodation, meaning, and implicature: Interdisciplinary foundations for pragmatics. In Philip R. Cohen, Jerry Morgan & Martha Pollack (eds.), *Intentions in communication*, 326–363. Cambridge, Massachusetts: MIT Press.
- Thomason, Richmond H. 2000. Modeling the beliefs of other agents. In Jack Minker (ed.), *Logic-based artificial intelligence*, 375–473. Dordrecht: Kluwer Academic Publishers.
- Traum, David R. & James F. Allen. 1994. Discourse obligations in dialogue processing. In *Thirty-second meeting of the Association for Computational Linguistics*, 1–8.
- Wooldridge, Michael J. 2000. *Reasoning about rational agents*. Cambridge, England: Cambridge University Press.

Richmond H. Thomason
Department of Philosophy
University of Michigan
Ann Arbor, MI 48109-1003
rthomaso@umich.edu

Matthew Stone
Department of Computer Science
Rutgers University
110 Frelinghuysen Road
Piscataway NJ 08854-8019
matthew.stone@rutgers.edu

David DeVault
USC Institute for Creative Technologies
12015 Waterfront Drive
Playa Vista, CA 90094-2536
devault@ict.usc.edu