

Coordinating Understanding and Generation in an Abductive Approach to Interpretation

Matthew Stone

Department of Computer Science
Rutgers University
mdstone@cs.rutgers.edu

Richmond H. Thomason

Philosophy Department
University of Michigan
rich@thomason.org

Abstract

We use a dynamic, context-sensitive approach to abductive interpretation to describe coordinated processes of understanding, generation and accommodation in dialogue. The agent updates the dialogue uniformly for its own and its interlocutors' utterances, by accommodating a new context, inferred abductively, in which utterance content is both true and prominent. The generator plans natural and comprehensible utterances by exploiting the same abductive preferences used in understanding. We illustrate our approach by formalizing and implementing some interactions between information structure and the form of referring expressions.

1 Introduction

The idea of interpretation as abduction is explained in (Hobbs et al., 1993) in terms of the following recipe for the interpretation of a sentence:

Prove the logical form of this sentence, together with the constraints that predicates impose on their arguments, allowing for coercions, merging redundancies where possible, making assumptions where necessary.

In (Stone and Thomason, 2002), we modify and extend this idea. We use a modal logic of context

to represent the logical form of utterances in terms of their potential to change the context. Accordingly, our abductive interpretations are proofs that describe a new context created by an utterance, in which its content is both true and prominent. Our earlier paper shows that the extension is both essential and effective in deriving correct abductive interpretations for sequences of utterances.

Here we go further to show how our approach supports the reasoning of a full conversational agent, capable of generation as well as understanding. Indeed, the approach reveals systematic commonalities in the reasoning required for understanding, generation and dialogue management—commonalities that spring from their use of shared representations and of abduction as the core reasoning process.

We illustrate the approach with an example in which the generation of appropriate information depends on attentional information, as well as on linguistic information and real-world knowledge.

2 Framework for Dialogue

We base our discussion on an information-state framework for dialogue management of the sort described, for instance, in (Larsson and Traum, 2000). Specifically, we assume that the information state takes the form of a comprehensive representation *cgn* of the common ground at state *n*, including not only the dynamic aspects of the dialogue but also the grammar and ontology that are consulted in proving interpretations. Since the information state will be used as a resource by inductive processes of understanding, the common

ground may consist not only of propositional information, but of attentional information that induces preferences between interpretations.

Participation in an ongoing dialogue is maintained by an update operation that revises the common ground in response to an utterance interpretation i , and a selection operation that draws on the agent’s private knowledge k as well as the common ground cgn to select a goal for generation. Understanding and generation mediate between utterances and communicative goals, and thereby construct the representations of interpretation over which dialogue update is defined.

Overall, the dialogue agent acts as a turn-taking manager, conforming to the following schematic perception-deliberation-action loop, where u, u' are utterances and i, i' are dynamic interpretations.

```
loop {
  input( $u$ );
  understand( $u, cgn, i$ );
   $cgn \leftarrow$  update( $i, cgn$ );
  generate(select( $k, cgn$ ),  $cgn, u', i'$ );
  output( $u'$ );
   $cgn \leftarrow$  update( $i', cgn$ ) }
```

This schema clearly shows the representational constraints that are imposed on interpretation and generation by the need for a single common ground update operation. The interpreter must produce the same update-supporting representations that are produced by the generator.

Constraints on the reasoning processes themselves that are used in understanding and generation follow from the further assumption that the achievement of mutuality in dialogue depends on similarities between the two processes. Suppose that a dialogue agent performs the step

generate(select(k, cgn), cgn, u', i'),

associating utterance u' with interpretation i' . It is natural to suppose that the same agent would also produce

understand(u', cgn, i');

that is, its interpretive component should derive the same effect that the generator intends. On this assumption, we can use similarities between dialogue agents to explain the maintenance of coordination in the course of a dialogue.¹

¹Since mistakes can be made about the discourse context, this model needs to be supplemented with an account of how miscoordinations can be prevented, identified and repaired.

3 Representing interpretation

To conform to this framework for dialogue, an abductive approach must provide parallel accounts of generation and understanding, in which the processes construct the same abductive interpretations despite the different goals and premises they use in reasoning. We now outline such an account, building on (Stone and Thomason, 2002), and illustrate it with the example utterance ‘He left’.

We continue to work with meaning representations formulated in a modal extension of Prolog. Interpretive goals for abductive proof are modalized atomic formulas and the clauses are modalized Prolog clauses whose component formulas may themselves be modalized. Modal operators represent contexts, which incorporate attentional as well as informational components.

We continue to assume that asserting a proposition has two effects; (i) the purely assertional effect of adding the proposition to the common ground suppositions of the conversation,² and (ii) the side effect of changing attentional features of the context. Developing this theme further, we use the formula *add-to-cg*(P, c_1, c_2) to say that adding the proposition P to the common ground and making the required attentional changes to c_1 will produce the new context c_2 .

Both understanding and generation seek to link an utterance with an intended change to the context. But in generation the content of the change is given as the goal of the plan, and what needs to be assumed is the utterance, as the means of achieving the goal. In understanding, what is given is the means of achieving a change, and what needs to be assumed is the intended content of this change.

Consider the utterance ‘He left’. The generator wishes to assert a proposition P , represented, say, by ‘*leave*(P, X)’, a combination of a predicate, ‘*leave*’, with a variable ‘ X ’ specifying a particular domain referent.³

The generation process in our example begins with the goal of deriving $c_1 : \text{add-to-cg}(P, c_1, c_2)$ from $c_1 : \text{leave}(P, X)$ —in the context c_1 the generator postulates a language-neutral description of

²See, for instance, (Stalnaker, 1981).

³In order to keep this example as simple as possible, we are ignoring all considerations having to do with the past tense.

(3.1)	$c_1 : he(X)$ $c_1 : utter('he left', E, P)$ $c_1 : do(E)$ $c_1 : leave(P, X)$ <hr style="width: 30%; margin-left: 0;"/> $c_1 : add-to-cg(P, c_1, c_2)$	Assumed, with low cost if X is masculine and in focus. Proved using the grammar, which selects the utterance as a way of expressing P . An intention is formed by hypothesizing an action: this is a low-cost assumption. Postulated in formulating the goal. Inference about the effects of communication.
(3.2)	$c_1 : he(X)$ $c_1 : utter('he left', E, P)$ $c_1 : do(E)$ $c_1 : leave(P, X)$ <hr style="width: 30%; margin-left: 0;"/> $c_1 : add-to-cg(P, c_1, c_2)$	Assumed, with low cost if X is masculine and in focus. Postulated based on observation. Postulated based on observation. Proved from utterance-type using grammar. Inference about the effects of communication.

a proposition, and attempts to use grammatical and contextual resources to find least-cost assumptions that allow the conversational goal to be proved. The generator is free to make assumptions that hypothesize the occurrence of a new utterance and describe its intended interpretation. As a side effect of our example proof, an intention is formed to utter ‘He left’.

Schematically, the proof has the structure of the abductive derivation shown in Proof (3.1). This proof uses shared information to explain how a proposed action or series of actions can update the context to assert P . Such a proof constitutes the generator’s discourse plan.

Understanding derives this same schematic proof, but by a different strategy. Understanding is given an eventuality E (the utterance), and the words that are uttered. This utterance needs to be classified as communicating a certain content. In this example, the process begins by postulating $c_1 : utter('he left', E, P)$ and adding it to the database. The goals of understanding and generation are exactly the same: proving that asserting the proposition P in c_1 will yield a new context c_2 . In this case, however, the new assumption that is added by the proof explains what P is.

When generation and understanding act reciprocally, generation’s intended P matches understanding’s inferred P and Proofs (3.1) and (3.2) give a common interpretation to both interlocu-

tors. Looking ahead to Section 6, we can adopt more fine-grained representations of context-dependence and context change in these interpretive proofs, so that an interlocutor can update the dialogue context simply by executing the transition specified there. In general, contexts incorporate different knowledge resources for conversational reasoning—we have already mentioned our division into informational resources (represented as axioms and rules) and attentional resources (represented as abductive preferences). We divide informational resources into old information and new information, which here has to do with new utterance events. Where c_n is a context, let i_n be the component of c_n representing old information, e_n be the component representing new information about events, and a_n be the attentional component. We can now revise Proofs (3.1) and (3.2) to make explicit the dependence of the utterance on these components of context, and the potential of the utterance to change them:

(3.3)	$a_1 : he(X)$ $e_1 : utter('he left', E, P)$ $e_1 : do(E)$ $i_1 : leave(P, X)$ <hr style="width: 30%; margin-left: 0;"/> $c_1 : add-info(P, i_1, i_2)$ $c_1 : put-in-focus(X, a_1, a_2)$
-------	---

The interpretation now provides two instructions

for the dialogue manager: it should update the informational context from state i_1 to state i_2 by assuming the proposition P that X left; and it should update the attentional state from a_1 to a_2 so that X remains a prominent potential referent for a pronoun.⁴ We will refer to this proof in our discussion of the implementation in Section 5.

4 Coordination

We now turn to the problem of coordination. Under what conditions can we in fact expect understanding to construct the interpretation that is intended by generation? And how easy will it be for understanding to do so? Concretely, consider examples (3.1) and (3.2), where generation produces an utterance U by deriving *add-to-cg*(P) from *leave*(P, X), and understanding derives *add-to-cg*(P) assuming the occurrence of U . Under what conditions will we expect understanding’s proof to precipitate the assumption *leave*(P, X) intended by the generator? And what inference will be involved?

4.1 Mutuality of information

We believe that coordination depends crucially on separating mutual and private information, as many dialogue architectures do, including the information-state architecture. In fact, this provides an important motivation to extend (Hobbs et al., 1993) along the lines we propose in (Stone and Thomason, 2002). Otherwise, the correct resolution of almost every utterance, including those discussed in (Hobbs et al., 1993), would depend on *ad hoc* assumptions about what is left out of the database, and/or *ad hoc* assignments of costs to axioms. In ‘He left’, for example, what if understanding knows about many people who could serve as the referent for ‘He’? Then its knowledge base will license an assumption *he*(Y) for many individuals Y . Abductive understanding will not infer that ‘He’ is X unless its preferences for this interpretation outweigh all the alternatives. These preferences must vary with circumstances.

In (Stone and Thomason, 2002), we propose to handle such effects in a general way by a straight-

⁴Thus we have $c_1 : \textit{put-in-focus}(X, a_1, a_2)$ in place of $a_1[X]a_2$ from (Stone and Thomason, 2002), and $a_1 : \textit{he}(X)$ in place of $a_1 : \textit{in-focus}^*(X)$ and $i_1 : \textit{man}(X)$.

forward modification of the scheme for abductive weights of (Hobbs et al., 1993). Discourse contexts specify the abductive weights that attach to assumptions. After an utterance where X has been set up as the most prominent referent (as a sentence subject, for example), the weights induce a low cost for assuming *he*(X) in the next utterance. Such specifications provide a very general approach to focus of attention; assumptions with relatively high costs become invisible, and different priorities can be assigned to the assumptions that are visible.

The reasoning needs of generation reveal another side to this requirement of mutuality. Generation starts from a specific communicative goal, but in deriving its discourse plan, it must respect the fact that this goal is a privileged, private resource. The generator must use the goal to guide the planning process, but not to constrain its reasoning about abductive interpretation. Otherwise, consider what would happen in contexts with many people who could serve as the referent for ‘He’. In building candidate interpretations, the generator would already know from its communicative goal that it wanted to produce an utterance whose subject referred to X . So to the generator—but not to another agent—it would look as if there were no alternative interpretations.

Our symmetric representations make it easy to maintain the dual perspective required in generation. The generator maintains multiple copies of interpretive proofs, which share the same structure. A distinguished proof records the instantiation required to establish the intended goal in generation; other proofs record the alternative instantiations derived only from shared information. The generator succeeds when only the distinguished proof remains, ensuring that the natural shared interpretation will achieve the goal. This dual representation extends the insights of the SPUD generator to abductive interpretation; see (Stone, 1998; Stone et al., 2003).

4.2 Coordination and preferences

Coordination in generation involves being clear, as well as being comprehensible. In our example, there will be many ways to identify X . Suppose it’s grounded that X is not only ‘He’, but also

‘The conference’s second presenter’. The generator hardly seems cooperative if it asks the understander to infer who ‘The conference’s second presenter’ is, when ‘He’ would do.

Again, the symmetric representations we adopt make it easy for the generator to take into account the interpretive effort required in understanding. In understanding, the cost of abductive proofs controls the search for possible interpretations. To keep this search space small, the generator should formulate an interpretation with a low abductive cost. Since the generator maintains its intended interpretation in the same representation it expects understanding to reconstruct, the generator can use the cost of this proof to guide its search. When the generator must make a choice between alternative expressions that could achieve its communicative goal, it can choose the one with the lowest cost. More generally, the generator can factor the cost of interpretation into its heuristics for searching the space of possible utterances.⁵ Concretely in our scenario, as long as the reference to X as ‘He’ has lower cost than a reference to X as ‘The conference’s second presenter’ (which it certainly will), the generator will prefer it.

The heuristics that underlie choice in generation architectures are often weakly motivated. That’s certainly true of SPUD. Our use of interpretive preferences draws on a collaborative view of dialogue to suggest a more principled alternative. And the approach offers a new perspective on the abductive weights themselves. They represent preferences over pairings of forms and meanings, which are used declaratively and reversibly throughout the dialogue architecture.

5 Implementation

We have implemented our model through a series of simple dialogue agents in Prolog. The rather limited ability of these agents to deliberate about

⁵We don’t pursue here the question of whether the generator must model the interpreter’s abductive cost, or can coordinate directly based on its own costs. The simplest and most direct way to achieve coordination would be for the generator to assume an interpreter that is like it in important respects. But this assumption will work only if dialogue agents can attune themselves to one another in ways that crucially affect the processes of generation and interpretation, and if the reasoning architecture treats these processes symmetrically.

appropriate conversational goals could be elaborated in a specific conversational domain. In addition, our current agents abstract away from surface realization and parsing; the utterances they exchange are represented as syntactic derivation trees rather than strings of words. It would be feasible to relax this simplification too. Nevertheless, these agents suffice as a testbed for exploring the predictions and opportunities of our architecture. For example, they realize the context-dependent patterns of reference generation and reference resolution described in Section 6.

Our implementation incorporates the abductive theorem prover described in (Stone and Thomason, 2002). It implements a modal extension of Prolog, and records its assumptions by manipulating *marked queries* consisting of a query $c : A$ and a label classifying the query as resolved, assumed or unsolved. An abductive proof is completed when all of its steps are either resolved or assumed. The reasoner prefers minimal-cost proofs, where cost is computed in terms of *abductive weights* that are attached to possible assumptions. The incorporation of abductive weights in contexts, formalized as modal operators, permits localized cost distributions to be used in proofs. A query $c : A$ is proved using the assumption costs incorporated in c .

The dialogue manager implements the loop described in Section 2. It includes several types of context updates, corresponding roughly to “speech acts”, which affect representations of the discourse state incorporating the notion of a *question under discussion* (see (Ginzburg, 1996)). At the same time, the current context is routinely revised in the course of a dialogue. Revision involves a mechanism for updating abductive costs using abstract, qualitative representations of preference, as in (Stone and Thomason, 2002).

Understanding starts from the syntactic structure of the utterance and the current context. It accesses the grammar to compute possible logical forms, and proves them abductively. Its interpretation is the overall least-cost proof. Conversely, generation implements a best-first search over grammatical derivations and assesses its progress based on the same model of utterance interpretation. As described in Section 4, the generator con-

structs the lowest-cost interpretation that will be understood as intended.

6 An Example

We will illustrate our approach with correspondences between informational state and the form of English referring expressions that are described in (Gundel et al., 1993). (For a related discussion of identifiability and activation, see (Lambrecht, 1994)[Chapter 3].) This paper postulates a “givenness hierarchy” with six different levels; here we will consider only the following part of the hierarchy: *IF* (*in focus*), *Act* (*activated*), *UI* (*uniquely identifiable*), and *TI* (*type identifiable*).

The levels of this hierarchy, in the order in which they were presented, correspond to the progressively weaker prominence of a reference. A referent that is in focus is not only known and actively under consideration, but because of the recent discourse or the mutual environment it is maximally prominent. A referent that is activated is in the current “short-term memory” of the conversation; it is readily available for retrieval. A referent is uniquely identifiable if there is an easy way of constructing a predicate that applies to the referent, and that distinguishes it for the hearer from all other referents. Finally, we treat type identifiability as a residual category.⁶ Note that we depart from (Gundel et al., 1993) in that their formulations are characterized in terms of the hearer’s knowledge. In view of the importance of mutuality, we substitute ‘grounded’ for ‘known by the hearer’.

Gundel et al. do not provide a detailed model of the cognitive state of a conversational agent. To incorporate their hierarchy in a discourse agent, we need to be more explicit. For current purposes, we can assume that the cognitive state consists of the following components:

⁶An item is type identifiable according to Gundel et al. if the hearer “is able to access a representation of the type of object described by the expression.” This last characterization strikes us as problematic. For one thing, we believe that the hierarchy is best thought of as applying to (discourse) referents, prior to referring expression generation. But in this case, it is not clear what would be meant by “the expression.” If you amend the characterization to require that the hearer should be able to access some predicate that in fact (whether or not the hearer knows it) applies to the item, then the classification applies to everything if ‘thing’ is a predicate.

- (i) A partially ordered set $\langle F, \preceq \rangle$ of items that are in focus (in conversational short-term memory and at the center of attention). The ordering \preceq represents relative prominence.
- (ii) A set $STM \supseteq F$ of items in short-term memory.
- (iii) A set $D \supseteq STM$ of referential distractors, and
- (iv) A set CG of common-ground predications, where a common-ground predication is a pair $\langle P, d \rangle$ consisting of a predicate P and a distractor d in D .
- (v) A set $DR \supseteq STM$ of discourse referents.

Our four givenness categories can then be defined in terms of cognitive state as follows:

- (1) $IF = F$.
- (2) $Act = STM$.
- (3) $UI = \{i : i \in D \text{ and there is a conjunction of propositions } P_1(d_1) \wedge \dots \wedge P_n(d_n), \text{ where each proposition is in } CG, \text{ such that } i \text{ is the only member of } D \text{ satisfying this conjunction}\}$.
- (4) $TI = DR$.

For simplicity, we assume that *IF* is a unit set.

Say that an item *is classified* by givenness category X if it belongs to X and to no more restrictive definiteness category. According to Gundel et. al, the pronouns ‘he’, ‘she’, and ‘it’ are appropriate for items classified by *IF*; definite NP’s with determiner ‘that’ or ‘this’ are appropriate for items classified by *Act*; definite NP’s with determiner ‘the’ are appropriate for items classified by *DR*; and indefinite NP’s with determiner ‘a’ are appropriate for items classified by *TI*. The following examples illustrate the correspondences.

- (6.1) The neighbor’s dog is a neighborhood nuisance.
It kept me awake last night.
- (6.2) I’m going to talk to the neighbor with the terrier.
That dog kept me awake last night.
- (6.3) [A dog is heard, barking outside the window.]
That dog kept me awake last night.

(6.4) I'm beginning to regret moving into this place.

The neighbor's dog kept me awake last night.

(6.5) I'm short on sleep.

A dog kept me awake last night.

We accept these correspondences. We now show how our conversational agent realizes them.

Our agent uses inference about interpretation to maintain the elements (i-v) of its cognitive state. The grammar of referring expressions includes rules such as these that update cognitive status:

(6.6) A sentence whose subject is discourse referent i creates a cognitive state in which $IF = \{i\}$.

(6.7) A discourse referent is activated when it is mentioned, and this activation persists until the end of the conversational episode.

These rules determine the specifications for context change in utterance interpretations, as specified in Proof (3.3). The discourse manager executes these specifications as part of context update. Nonlinguistic events induce similar updates:

(6.8) A discourse referent is assigned to a new entity that is perceptually salient in the mutual perceptual environment, and is assigned an activated status.

Now, our agent exploits these connections, by once more reasoning about interpretation in context. Our grammar specifies preferred associations between cognitive state and linguistic patterns within a specific context. For instance, we link pronouns to referents by (6.9).

(6.9) A context in which $IF = \{i\}$ is associated with linguistic preferences for a pronominal NP when i is the reference of that NP,

In combination with (6.6), (6.9) will cause the interpreter to prefer X as the referent to the utterance of "He" in (6.1). The very same mechanism will cause the generator to prefer "He" here: among the recognizable references to X , this utterance has the lowest cost.

7 Discussion

The original theory of interpretation as abduction contributed uniform analyses for a wide range of pragmatic phenomena in isolated sentences. We wish to claim that our new approach offers a similar benefit for dialogue; it formalizes analyses in common terms that can be combined or reconciled with one another. However, readers familiar with the formalization of interpretation in (Hobbs et al., 1993) will note that proofs such as (3.3) are quite different from the interpretation format of that paper. These differences are natural consequences of rethinking interpretation so that it applies to utterances rather than merely to sentences, and to ensure that it takes only into account information that is shared. The format of Proof (3.3) also differs slightly from the account of interpretive proofs given in (Stone and Thomason, 2002); these differences have mainly to do with the need to provide a common form for interpretation and generation. The differences between Proof (3.1) and the account of the same example in (Stone, 2003) have mainly to do with the introduction of explicit speech acts, represented as operators on the common ground.

Individually, the principles of our account of English referring expressions are familiar. Interpretation involves inference that combines linguistic meaning and an extensive common ground: see e.g., (Clark and Marshall, 1981). Grounding a new utterance involves coordinated updates to the attentional state of the dialogue; see e.g., (Brennan, 1998). Interpretation means recovering the most preferred interpretation; see e.g., (Hobbs et al., 1993). Conversely, generation means using the most preferred form that can be understood in context; see e.g., (Buchwald et al., 2002). We regard it as a strength of our formalism that its key principles are not controversial.

Nevertheless, formalizing (6.1–6.5) as we have done crucially requires a framework in which all three principles are simultaneously available. No previous formalism does this. For example, treatments of reference in generation such as (Stone et al., 2003) model the information that interlocutors use to disambiguate expressions, but require additional mechanisms if they are to select expres-

sions with a suitable form. Conversely, treatments of pronouns in discourse such as (Buchwald et al., 2002) can determine whether a pronoun naturally evokes a target referent, but do not generalize to the construction of referring expressions when a pronoun would be ambiguous. Our proposal reconciles the insights of both approaches in a single computational formalism.

8 Conclusion

In this paper, we have presented an approach to the coordination of generation and interpretation in dialogue that adds a principled treatment of discourse context to the interpretation as abduction framework. Coordination depends on shared background that can inform interpretation and on shared preferences that say which interpretations are natural. In our architecture, generation and understanding both rely on these preferences to derive good interpretations. The architecture explains how speakers can use knowledge of language to produce and understand the concise and comprehensible utterances they seem naturally to use.

The flexibility of the approach gives new impetus to efforts to analyze extended natural dialogues in formal terms. We are optimistic that such projects will confirm the elegance and power not only of the formalism itself, but of the intuitive principles of coordination in dialogue that it embraces.

References

- Susan E. Brennan. 1998. Centering as a psychological resource for achieving joint reference in spontaneous discourse. In Marilyn A. Walker, Aravind K. Joshi, and Ellen Prince, editors, *Centering Theory in Discourse*, pages 227–249. Oxford University Press.
- Adam Buchwald, Oren Schwartz, Amanda Seidl, and Paul Smolensky. 2002. Recoverability optimality theory: Discourse anaphora in a bidirectional framework. In Johan Bos, Mary Ellen Foster, and Colin Matheson, editors, *EDIALOG 2002: Proceedings of the Sixth Workshop on the Semantics and Pragmatics of Dialogue*, pages 37–44. Cognitive Science Centre, University of Edinburgh, Edinburgh.
- Herbert H. Clark and Catherine R. Marshall. 1981. Definite reference and mutual knowledge. In Arivind Joshi, Bonnie Webber, and Ivan Sag, editors, *Elements of Discourse Understanding*, pages 10–63. Cambridge University Press, Cambridge, England.
- Jonathan Ginzburg. 1996. Interrogatives: Questions, facts, and dialogue. In Shalom Lappin, editor, *The Handbook of Contemporary Semantic Theory*, pages 385–422. Blackwell Publishers, Oxford.
- Jeanette K. Gundel, Nancy Hedberg, and Ron Zacharski. 1993. Cognitive status and the form of referring expressions in discourse. *Language*, 69(2):274–307.
- Jerry Hobbs, Mark Stickel, Douglas Appelt, and Paul Martin. 1993. Interpretation as abduction. *Artificial Intelligence*, 63(1–2):69–142.
- Knud L. Lambrecht. 1994. *Information Structure and Sentence Form: Topic, Focus, and the Mental Representations of Discourse Referents*. Cambridge University Press, Cambridge, England.
- Staffan Larsson and David Traum. 2000. Information state and dialogue management in the TRINDI Dialogue Move Engine Toolkit. *Natural Language Engineering*, 6:323–340.
- Robert C. Stalnaker. 1981. Assertion. In Peter Cole, editor, *Radical Pragmatics*. Academic Press, New York.
- Matthew Stone and Richmond H. Thomason. 2002. Context in abductive interpretation. In Johan Bos, Mary Ellen Foster, and Colin Matheson, editors, *EDIALOG 2002: Proceedings of the Sixth Workshop on the Semantics and Pragmatics of Dialogue*, pages 169–176. Cognitive Science Centre, University of Edinburgh, Edinburgh.
- Matthew Stone, Christine Doran, Bonnie Webber, Tonia Bleam, and Martha Palmer. 2003. Microplanning with communicative intentions: The SPUD system. *Computational Intelligence*, 19(4). To appear.
- Matthew Stone. 1998. *Modality in Dialogue: Planning, Pragmatics and Computation*. Ph.D. dissertation, Computer Science Department, University of Pennsylvania, Philadelphia, Pennsylvania.
- Matthew Stone. 2003. Communicative intentions and conversational processes in human-human and human-computer dialogue. In John Trueswell and Michael Tanenhaus, editors, *World Situated Language Use*. The MIT Press, Cambridge, Massachusetts. To appear.