

Incorporating Asymmetric and Asynchronous Evidence of Understanding in a Grounding Model

Alexandre Denis, Guillaume Pitel, Matthieu Quignard, Patrick Blackburn

TALARIS team

UMR 7503 LORIA/INRIA Lorraine

Campus scientifique, BP 239

F – 54506 Vandoeuvre-lès-Nancy cedex

alexandre.denis@loria.fr

Abstract

The grounding process relies on the evidence that speakers give about their understanding (Clark and Schaefer, 1989). However in existing formal models of grounding (Cahn, 1992; Cahn and Brennan, 1999; Traum, 1999) evidence of understanding is assumed to be *symmetrically* and *synchronously* shared by the speakers. We propose a formal model, based on (Cahn, 1992), that removes these simplifications; we do so by distinguishing the phase of interpretation from the phase of evidence extraction and introducing the notion of floating contributions.

1 Introduction

A dialogue is a process that presupposes the collaboration of both participants. Each speaker in turn assumes that the other will show evidence of understanding or misunderstanding of her utterance, and evidence indicating its relevance to the previous utterance. This mutual assumption is the basis of the *grounding process* (Clark and Wilkes-Gibbs, 1986; Clark and Schaefer, 1989), the process by which speakers try to reach mutual understanding. Successful grounding does not guarantee mutual understanding though: it can happen that the grounding evidence leads two speakers to believe that they have achieved perfect understanding, whereas in reality they have understood two completely different things (Cherubini and van der Pol, 2005). But although this shows that successful grounding is not a *sufficient* condition for achieving mutual understanding, it does seem to be a *necessary* one.

Different models of the grounding process define when (and, sometimes, how) an utterance is

added to the common ground (a representation of what is believed to have been mutually accepted). In the *Contribution Model* (Clark and Schaefer, 1989) the grounding process results in a recursively structured directed acyclic graph representation of the dialogue grounding structure, the basic unit of which is the *contribution*. Contributions are twofold units consisting of: (1) an utterance called the *presentation* (or *Pr*) and (2) an *acceptance* linked to a sequence of contributions or a single utterance. The acceptance (or *Ac*) contains the negotiation of the understanding of the presentation in order to reach the *grounding criterion*. The grounding criterion is a threshold defined by Clark and Wilkes-Gibbs (1986) to represent the level of understanding required by the contributor; we shall use the expression *grounding status* to mean the current state of the believed mutual understanding of an utterance. When the grounding criterion holds for a contribution, that is, when its status is grounded, both speakers consider it closed and can choose whether or not to integrate its semantic content as a mutual belief. The grounding status is established via simple evidence of understanding and relevance. The Contribution Model was the pioneering approach to the modeling of grounding and its insights influenced the subsequent development of formal models intended for computational applications.

Probably the best known of these subsequent models is the *Grounding Acts Model* (Traum and Allen, 1992; Traum, 1994; Traum, 1999). This model is based on the notion of *grounding acts*, low level communicative acts whose goal is to ground content at the utterance level. The basic unit of analysis provided by the Grounding Acts Model is a non-recursive sequence of utterances called a *Discourse Unit* (DU). The grounding pro-

cess is modelled by an update of the state of a Discourse Unit by a grounding act; this makes the approach particularly suitable for integration into information-state based models of dialogue (such as Matheson et al. (2000)); transitions between states are modelled in (Traum and Allen, 1992) and many subsequent papers using finite state automata triggered by the various grounding acts. For example, a *RequestRepair* act by participant A would send a Discourse Unit into a state where a *Repair* by participant B and a subsequent *Acceptance* by A would be needed to ground it.

The Grounding Acts Model makes the assumption that the grounding level can be distinguished from the intentional level. However, as was noted by (Stirling et al., 2000), it is often not easy to delineate DUs, which makes it difficult to clearly distinguish the grounding level from deeper levels of understanding that emerge via complex exchanges. Hence, as our primary motivation is to explore ways of uniformly integrating grounding at the utterance level with complex negotiations of understanding, we have not taken the Grounding Acts Model as our point of departure.

Instead we have chosen to develop the *Exchange Model* approach presented in (Cahn, 1992; Cahn and Brennan, 1999), which are more directly based on the original Contribution Model. The central innovation provided by Exchange Models is a level of *exchange* that is higher than the level of contributions (this central notion is very much in the spirit of the implicit adjacency pairs used in Clark and Schaefer (1989)). Like work based on the Grounding Acts Model, these Exchange Models have a formal definition and provide on-line models of grounding. What makes them particularly useful for our purposes, however, is that they follow the Contribution Model in producing graph-like representations of the dialogue grounding structure; in our view, this makes them particularly well-suited for modeling more complex negotiations of understanding.

Nonetheless, different as these three types of model are, they share a common deficiency: *they cannot deal with wrongly recognized or unrecognized grounding acts or evidence of understanding*. In the original Exchange Models, the evidence is always assumed to be *symmetric* and *synchronous*—that is *correctly* and *immediately* understood by the hearer. In the Grounding Acts Model, matters are a little more subtle. There an

unrecognized grounding act would initiate a new Discourse Unit, and hence the model might be said to handle asymmetric grounding. Nonetheless, it is not obvious how, once grounded, this Discourse Unit should be reintegrated, nor how the effects of the newly understood grounding act could be taken into account with respect to previous Discourse Units. It may be the case that the Grounding Acts Model and related information states approaches (such as Matheson et al. (2000; Larsson and Traum (2000)) could be extended to handle this kind of reintegration, perhaps by providing additional update rules. But we have found that the (recursive) graph-like representations used by Exchange Models provides a particularly perspicuous setting for a preliminary explorations of the issues involved.

Accordingly, we shall proceed as follows. In Section 2 we discuss existing Exchange Models and their shortcomings in more detail. In Section 3, we present an augmented Exchange Model, inspired by (Cahn, 1992), which repairs these deficiencies. In Section 4, with the help of an example, we show in detail how the model works. Section 5 concludes.

2 A closer look at Exchange Models

The Exchange Models proposed in (Cahn, 1992; Cahn and Brennan, 1999) are intended to formalize the Contribution Model to enable it to be embedded in dialogue systems. Like the Contribution Model, they are based on (recursive) graph-like structures, but they add a level of *exchange* above the Contribution/Presentation/Acceptance levels present in the Contribution Model. An exchange is a pair of contributions defined relative to a task: the first contribution proposes a task while the second contribution executes the task. The grounding process itself is modeled by a decision table based on two features: (1) the evidence of understanding manifested in an utterance and (2) the role of the current utterance in an exchange (i.e. in a dialogue task). In these models, the grounded dialogue structure is represented from each speakers' individual point of view, and "*all contribution graphs are private models, and can represent the perspective of only one agent*" (Cahn and Brennan, 1999).

The model defined in (Cahn, 1992) (henceforth *EM92*) uses three categories of evidence: *UNDERSTOODRELEVANT*, *NOTUNDERSTOOD* and

u_1 : Where does Dan work ?
 s_2 : In the natural language group
 u_3 : No, I meant his office

Figure 1: UNDERSTOODNOTRELEVANT evidence example (dialogue 6.3.3)

UNDERSTOODNOTRELEVANT. However it considers the integration of the utterances of the speaker based on the evidence she *intended* to produce, and does not take into account *how the hearer actually interprets this evidence*. This model can thus only render the perspective of the speaker. The model defined in (Cahn and Brennan, 1999) (henceforth *EM99*) is limited to the system’s point of view, unlike *EM92* which is participant-agnostic. It is based on two categories of evidence: ACCEPTABLE and NOTACCEPTABLE for the user. This is not enough to cover all cases; for example, it is not possible for the system to warn that its utterance was misunderstood (though the user can do this). The main improvement of *EM99* over *EM92* is that it characterizes grounding from the hearer’s point of view (in this case the system) and integrates the utterances of the user as well as those of the system.

So what is wrong with these models? An example should make matters clear. Consider Figure 1, which shows an example dialogue from (Cahn, 1992). When S utters s_2 , she believes her utterance is a relevant answer to u_1 . That is, she appends s_2 as the second contribution of the exchange initiated by u_1 (see Figure 2).¹




Figure 2: Dialogue 6.3.3 after s_2

When S receives u_3 and extracts the UNDERSTOODNOTRELEVANT evidence it conveys, she has to restructure her view of the dialogue to take into account the fact that s_2 is *not* relevant with respect to u_1 . That is, she must move s_2 into a new exchange in the acceptance phase of u_1 and append u_3 as the second contribution of this exchange (see Figure 3).

¹In the paper we sometimes use the term utterance instead of contribution, or an utterance symbol to denote a contribution; in these case we always mean “the contribution presented by this utterance”.

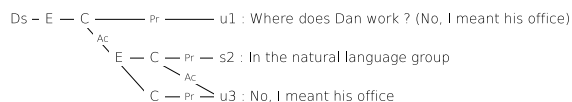


Figure 3: Dialogue 6.3.3 after u_3

But now take the point of view of U on the situation: she interprets s_2 as not relevant with respect to u_1 *as soon as she receives it*, and thus does *not* have to revise her model. Instead she integrates s_2 directly into the acceptance phase of her own utterance u_1 . To properly model such situations from both points of view, we need to consider not merely the intended evidence of understanding, but also the actual interpretation of evidence.

However *EM92* cannot handle this behavior because it does not consider the interpretation of the hearer: it is focused on how integrate an utterance according to the evidence of understanding the speaker intended it to convey, and not on how the hearer actually interpreted the utterance. The model *EM99*, on the other hand, partially takes into account the hearer’s interpretation in one case, namely when the user gives ACCEPTABLE evidence and does not propose a task. For example, consider the exchange shown in Figure 4. *EM99* handles such examples well. If the system believes it does not understand u_2 , its next utterance would initiate a new exchange in the acceptance phase of u_2 . On the other hand, if it believes it understands u_2 , its next utterance would initiate a new exchange at the dialogue level.

s_1 : Where’s Dan ?
 u_2 : In his office (+ noise)

Figure 4: ACCEPTABLE evidence example

This is a step in the right direction, but it does not go far enough: the interpretation of the hearer needs to be taken into account in *all* cases, for *both* speakers, *whatever* the evidence conveyed by the interpreted utterance. Only if the hearer understands can she extract the evidence of understanding. If the interpretation of the hearer is not taken into account, it is impossible to consider the non-understanding (or the misunderstanding) of the evidence of understanding. As a consequence, models that make this simplification can only deal with *symmetric* and *synchronous* evi-

u_1 : Where's Dan ?
 s_2 : Dan Smith or Dan Jones ? (+ noise)
 u_3 : Uh, what did you say ?

Figure 5: Not understanding a NOTUNDERSTOOD evidence of understanding

dence of understanding. Symmetric means that the presented evidence of understanding is always understood by the hearer as expected by the speaker. Synchronous means that it is also understood as soon as the utterance is emitted. Neither *EM92* nor *EM99* can handle the fact that the *acceptance function* of an utterance is not always played at the same moment for the two speakers. In the example in Figure 5, the user does not know that the utterance s_2 manifests NOTUNDERSTOOD evidence because she herself has not understood s_2 sufficiently to extract the evidence of understanding.

3 An extended Exchange Model

The aim of our work is to specify a grounding model that handles asymmetric and asynchronous evidence of understanding. We do so in a way that retains the advantages of both *EM92* and *EM99*, and follow these models in using a small collection of sharp understanding categories. We found *EM92*, which does not distinguish the user from the system and has better categories of understanding, to be a better starting point, and thus have reworked the key insights of *EM99* in the setting of *EM92*. We introduce asymmetry and asynchronicity into this model by distinguishing two steps in the interpretative process: the understanding of an utterance, and the extraction of the evidence of understanding it conveys. Our model is defined from the point of view of the hearer (who we will refer to by *self*). It considers how *self* understands an utterance *before* taking into account the evidence of understanding it shows. If the utterance is UNDERSTOODRELEVANT, the evidence it shows is extracted and is integrated as in *EM99*. If the utterance is UNDERSTOODNOTRELEVANT, it is integrated as initiating a new exchange under the acceptance phase of the previous contribution without considering the evidence of understanding it shows. If the utterance is NOTUNDERSTOOD, the contribution is introduced as *floating*, waiting for later integration in the graph.² Its integration into

²In this version, NOTUNDERSTOOD means that the evidence is not understood either. Multiple degrees of under-

the main dialogue structure is possible only when its acceptance phase shows what the evidence of understanding was. Then the utterance has to be reinterpreted, taking into account the newly understood evidence. However the grounding status of floating items could remain pending and never be solved, for example if the acceptance phase is abandoned.

Do floating contributions have an analog in the Grounding Acts Model? We don't believe so. One could try comparing the collection of ungrounded states of a Discourse Unit to the acceptance phase of a contribution. That is, the open state of the acceptance phase of a contribution in an Exchange Model could be regarded as the analog of the ungrounded states of a Discourse Unit in the Grounding Acts Model. But the notion of a floating contribution is stronger than the notion of ungrounded states: when a contribution is floating, it means not only that it is not grounded yet, but also that the evidence it manifests is not known either. Furthermore, floating status isn't correctly captured by the ungroundable state used in the Grounding Acts Model either. The ungroundable state is a terminal state, reached by canceling the grounding process. However the floating status that some contributions acquire in our approach is intended to be temporary—if the evidence conveyed by the contribution comes to be understood, its floating status is cancelled and the contribution is integrated into the main dialogue structure.

So: how can we augment the Exchange Model to handle asymmetric and asynchronous evidence of understanding? The main additions we shall make are the following. First, in order to keep track of the floating contributions, we have to maintain another structure which contains the sequence of pending contributions. Second, we have to handle reinterpretation and specify how the newly acquired evidence of understanding is used to integrate a floating contribution. Third, an utterance can now give evidence of understanding concerning many previous utterances, because an utterance which closes an acceptance phase and gives evidence as such can now reveal how to interpret the accepted contribution too.

A formalization³ is presented in Table 1 and Table 2. The tables are simplified. We do not discuss here dialogue beginning and ending nor the necessity of having a three-fold context $\langle S_i, O_j, S_k \rangle$ to manage reinterpretation. The actual implementation also deals with evaluation utter-

standing are possible though (Brennan and Hulsteen, 1995).

³The tables are simplified. We do not discuss here dialogue beginning and ending nor the necessity of having a three-fold context $\langle S_i, O_j, S_k \rangle$ to manage reinterpretation. The actual implementation also deals with evaluation utter-

<i>Self interpretation of O_i</i>	<i>Integration of O_i</i>	<i>Integration of S_{i+1}</i>
UNDERSTOODRELEVANT w/r S_{i-1}	integrate O_i according to the evidences of understanding it shows: if O_i shows evidences about another utterance S_j , call Table 2 with S_j else call it with S_{i-1}	integrate S_{i+1} after O_i
UNDERSTOODNOTRELEVANT w/r S_{i-1}	integrate O_i as initiating a new exchange in the acceptance phase of S_{i-1}	integrate S_{i+1} after O_i
NOTUNDERSTOOD	O_i presents a floating contribution, waiting to be understood to be integrated	integrate S_{i+1} as initiating a new exchange in the acceptance phase of O_i

Table 1: Integration of both utterances (O_i, S_{i+1})

ble 2. It covers the aforementioned cases with delays in the integration of utterances. We use the following notation: O_i stands for the utterance produced by the *other* speaker (O) at time i , and S_{i+1} stands for the utterance produced by *self* (S) at time $i + 1$. An utterance is said to initiate an exchange if it is the presentation of the first contribution of this exchange. An exchange is said to be open if its first contribution is set while its second contribution is not. The main dialogue structure is called D and the floating structure F . Finally, “integrate u_i after u_j ” is a shorthand for:

- if u_j initiated an exchange, append u_i as the second contribution of this exchange;
- else append u_i as the second contribution of the closest upper level open exchange, if there is one;
- else (all exchanges are closed), u_i initiates a new exchange at the dialogue level.
- in all cases u_i closes the acceptance phase of u_j .

Reinterpretation of an utterance consists of calling the algorithm again with a new interpretation and new evidence of understanding. The only difference is that the contribution presented by this utterance does not have to be created because it already exists in the floating structure. If the utterance is eventually understood (relevant or not) it can be moved in the dialogue structure in accordance with its new interpretation, and the new evidence of understanding it shows. This evidence of understanding is consequently acquired asynchronously by the two participants.

The main simplifying assumptions made by our algorithm are the following:

ances and abandons. For a complete description, please refer to <http://www.loria.fr/~denis>

- We suppose a direct correlation between the result of an interpretation of O_i and the evidence of understanding conveyed by S_{i+1} . How S interprets O_i is manifested in the utterance she produces in turn S_{i+1} . That is, if an utterance is not understood or not relevant, one has to clarify the situation. This simplification is based on the collaborative dialogue hypothesis.
- In the version of the algorithm presented above, the evidence of understanding is either understood or not. That is, the asymmetry is binary and there cannot be any misunderstanding of the evidence of understanding. Such misunderstandings would be more complex to handle because of the increased divergence between the participants dialogue representation structures. But systems can be mistaken when extracting the evidence of understanding, and we think it will be necessary for dialogue systems to represent this.
- Contributions always alternate. The present algorithm does not actually manage several contributions in one speech turn because this would mean taking into account *interleaved* evidence of understanding. But, once again, we feel that this extension will be necessary to handle more realistic dialogues.

A tool illustrating our model has been implemented in Java. This takes as input a dialogue where each utterance is annotated by the evidence of understanding its speaker believes it to convey. The resulting output is the dialogue structure and the floating structure for both speakers at different steps. The tool was used to generate the diagrams used in this paper, and in particular, the diagrams in the example to which we now turn.

Evidence of understanding of S_j showed in O_i	S_j did not initiate an exchange or initiated an exchange at the dialogue level	S_j initiated an exchange in an acceptance phase
UNDERSTOODRELEVANT w/r O_{j-1}	integrate O_i after S_j	integrate O_i after S_j , if the accepted contribution is flouting, reinterpret its presentation O_k : call Table 1 where $O_i = O_k$
UNDERSTOODNOTRELEVANT w/r O_{j-1}	move the S_j contribution as the first contribution of a new exchange in the acceptance phase of O_{j-1} , integrate O_i after S_j	
NOTUNDERSTOOD	integrate O_i as initiating a new exchange in the acceptance phase of S_j	

Table 2: Integration of O_i when it is understood and thought relevant

Utterance and evidence of understanding it shows	Point of view of A	Point of view of B
a_1 : Where does Dan work ?	Da - E - C - Pr - a1	Db - E - C - Pr - a1
b_2 : In the natural language group UNDERSTOODRELEVANT(a_1)	Da - E - C - Pr - a1 Fa - C - Pr - b2	Db - E - C - Pr - a1 C - Pr - b2 AC
a_3 : What did you say ? NOTUNDERSTOOD(b_2)	Da - E - C - Pr - a1 Fa - C - Pr - b2 AC E - C - Pr - a3	Db - E - C - Pr - a1 C - Pr - b2 AC E - C - Pr - a3
b_4 : I said: in the natural language group UNDERSTOODRELEVANT(a_3)	Da - E - C - Pr - a1 E - C - Pr - b2 AC E - C - Pr - a3 AC C - Pr - b4	Db - E - C - Pr - a1 C - Pr - b2 AC E - C - Pr - a3 AC C - Pr - b4
a_5 : No, I meant his office UNDERSTOODRELEVANT(b_4) UNDERSTOODNOTRELEVANT(b_2)	Da - E - C - Pr - a1 E - C - Pr - b2 AC E - C - Pr - a3 AC C - Pr - b4 AC C - Pr - a5	Db - E - C - Pr - a1 E - C - Pr - b2 AC E - C - Pr - a3 AC C - Pr - b4 AC C - Pr - a5
b_6 : Near post H33 UNDERSTOODRELEVANT(a_5)	Da - E - C - Pr - a1 E - C - Pr - b2 AC E - C - Pr - a3 AC C - Pr - b4 AC C - Pr - a5 AC C - Pr - b6	Db - E - C - Pr - a1 E - C - Pr - b2 AC E - C - Pr - a3 AC C - Pr - b4 AC C - Pr - a5 AC C - Pr - b6

Table 3: Detailed example

4 Detailed example

The example in Table 3 is a modification of the example 6.3.3 in (Cahn, 1992), in which the second utterance is not understood by the user (called *A* whereas the other participant, the system, is called *B*). This dialogue illustrates the asymmetry and asynchronicity of the learning of the evidence of understanding showed by b_2 . The left column presents the utterances and the evidence of understanding showed by them from their speaker's point of view. The two other columns present the dialogue structure according to each point of view.

The first utterance a_1 is believed UNDERSTOODRELEVANT⁴ by *B* and is integrated normally as initiating an exchange at the dialogue level. The dialogue viewed by *A* is the same.

The second utterance b_2 shows a divergence. *B* believes that b_2 presents an UNDERSTOODRELEVANT evidence of understanding and thus integrates it as the second contribution of the first exchange. However this evidence is not shared by *A*, who does not understand b_2 and therefore cannot integrate it. She just keeps the contribution floating, awaiting to be integrated when it is sufficiently understood (see F_a in Table 3).

Utterance a_3 shows that b_2 was NOTUNDERSTOOD by *A* and that she requires clarification. Because a_3 is understood by *B*, the evidence of understanding it contains is used to integrate it as the initiator of a new exchange under the acceptance phase of b_2 contribution.

Utterance b_4 shows that a_3 was interpreted UNDERSTOODRELEVANT by *B*. Therefore it is integrated by both speakers as the second contribution of the clarification exchange. However there is a new divergence when processing the utterance b_4 . For *B*, b_4 is only an answer to the clarification request. But with b_4 , *A* can now interpret b_2 . As b_2 is now understood by *A*, she can extract the evidence of understanding it showed, and act according to her own interpretation. In this case, because b_2 is UNDERSTOODNOTRELEVANT by *A*, she won't take into account the evidence of understanding showed by b_2 . The acquisition of the evidence of understanding showed by b_2 is asynchronous but not taken into account; see Table 1. The reinterpretation of b_2 , according to the UNDERSTOODNOTRELEVANT rule, leads to the b_2 contribution being embedded as initiating a new

⁴The first utterance is assumed relevant when it is understood

exchange under the acceptance phase of a_1 .

The utterance a_5 is crucial for reaching the grounding criterion. It makes available two pieces of evidence of understanding: first it shows that b_4 is an UNDERSTOODRELEVANT reply to a_3 and second it shows that b_2 is UNDERSTOODNOTRELEVANT with respect to a_1 . Its effect is to revise the *B* view on the dialogue to create a new exchange in the acceptance phase of a_1 . Doing this means that the structures of the grounding model converge for both speakers; they now agree on the current view of dialogue.

The utterance b_6 is the final answer to the first question. It shows that a_5 was UNDERSTOODRELEVANT by *B* and UNDERSTOODRELEVANT by *A*. It is integrated as a relevant reply to the first contribution of the upper level exchange.

5 Discussion and further work

This paper discusses the problems posed by asymmetric and asynchronous evidence of understanding, and gives a preliminary model of how such evidence could be handled. It does so by distinguishing the phase of interpretation from the phase of evidence extraction and introducing the notion of floating contributions into the Exchange Model. Such contributions cannot be immediately attached to the dialogue structure because the evidence of understanding they show is not known. When these contributions are accepted, they have to be reinterpreted in order to extract the evidence of understanding they manifest.

A side effect of our model is that it provides a novel solution to the recursive acceptance problem defined in (Traum, 1994; Traum, 1999): if an acceptance utterance needs to be accepted *before* it can play its acceptance function, then no contribution would ever be complete. To solve the problem, we make the assumption that a participant may form the belief that she has understood (or not) an utterance as soon as she receives it; she does *not* have to subordinate her belief to further acceptance (we believe that this assumption can be motivated by the ideas on timing in joint actions in Chapter 3 of (Clark, 1996)). The acceptance function of an utterance can be played, from the hearer's point of view, *as soon as she understands the utterance*. On the other hand, to check whether what she said successfully played its intended acceptance role, the speaker of the utterance has to wait for the hearer's response. However, *as soon*

as the hearer responds, the appropriate acceptance function may be played. But when misunderstanding occurs, the acceptance role of an utterance is delayed up to the moment it is sufficiently understood to be integrated into the common ground.

The implemented model we have presented still suffers from a number of limitations; for example it does not deal with misunderstanding of the evidence of understanding. Planned future work will cover these more complex divergences in dialogue structure in addition to multi-contributions, that is, when several contributions by the same speaker in the same turn. We hope that this model and its implementation will be the first stage of a larger enterprise: specifying the grounding status of the contents of a contribution in terms of dialogue structure.

Acknowledgements

We would like to thank TALARIS members Luciana Benotti and Daniel Coulon for feedback and discussion during the writing of this paper.

References

- Susan E. Brennan and Eric A. Hulteen. 1995. Interaction and feedback in a spoken language system: a theoretical framework. *Knowledge-Based Systems*, 8:143–151.
- Janet E. Cahn and Susan E. Brennan. 1999. A psychological model of grounding and repair in dialog. In *AAAI Fall Symposium on Psychological Models of Communication in Collaborative Systems*, pages 25–33.
- Janet E. Cahn. 1992. A computational architecture for the progression of mutual understanding in dialog. Technical Report 92-4, Music and Cognition Group M.I.T. Media Laboratory.
- Mauro Cherubini and Jakko van der Pol. 2005. Grounding is not shared understanding: Distinguishing grounding at an utterance and knowledge level. In *CONTEXT'05, the Fifth International and Interdisciplinary Conference on Modeling and Using Context*, Paris, France, July 5-8.
- Herbert H. Clark and Edward F. Schaefer. 1989. Contributing to discourse. *Cognitive Science*, 13:259–294.
- Herbert H. Clark and Deanna Wilkes-Gibbs. 1986. Referring as collaborative process. *Cognition*, 22:1–39.
- Herbert H. Clark. 1996. *Using Language*. Cambridge University Press.
- Staffan Larsson and David Traum. 2000. Information state and dialogue management in the trindi dialogue move engine toolkit. *Natural Language Engineering*, 6(3-4):323–340.
- Colin Matheson, Massimo Poesio, and David Traum. 2000. Modelling grounding and discourse obligations using update rules. In *Proceedings of NAACL 2000*, May.
- Lesley Stirling, Ilana Mushin, Janet Fletcher, and Roger Wales. 2000. The nature of common ground units: an empirical analysis using map task dialogues. In *Gotalog 2000, 4th Workshop on the Semantics and Pragmatics of Dialogue*, pages 159–165, Gothenburg, Sweden.
- David R. Traum and James F. Allen. 1992. A "speech acts" approach to grounding in conversation. In *International Conference on Spoken Language Processing*, pages 137–40.
- David R. Traum. 1994. *A Computational Theory of Grounding in Natural Language Conversation*. Ph.D. thesis, University of Rochester, Rochester, New York.
- David R. Traum. 1999. Computational models of grounding in collaborative systems. In *working notes of AAAI Fall Symposium on Psychological Models of Communication in Collaborative Systems*, pages 124–131.