An empirically based computational model of grounding in dialogue

Harry Bunt, Roser Morante and Simon Keizer
Department of Communication and Information Sciences
Tilburg University
{r.morante, harry.bunt,s.keizer@uvt.nl}

Abstract

In this paper we present a simple, empirically grounded computational model of grounding in dialogue. Grounding is shown to occur as a result of the dynamics of the information states of dialogue participants. A step-by-step analysis and representation of how information states develop through dialogue utterance processing illustrates exactly how this works.

1 Introduction

In an information-state update (ISU) approach, a dialogue is viewed as a sequential structure consisting of communicative acts that the participants perform in order to change each other's information state. For example, consider the following dialogue at a railway station between traveler A and employee B of the railway company:

(1) 1. A: Excuse me, can you tell me what time the next train to Amsterdam leaves?
2. B: Yes, that's at 9:17.
3. A: And at which platform is that?
4. B: That's at platform 5.
5. A: Thanks a lot.
6. B: You're welcome.

The second utterance tells A, among other things, that B believes that the next train to Amsterdam leaves at 9:17. Let us call this information p. Assuming that employees of the railway company provide correct information about train departure times, A will adopt the belief that p. So both participants now believe that p, and A also believes that B believes that p. After utterance 3, B will moreover believe that A has come to believe that p, although nothing is said about that. The dialogue continues on the topic of departure platform, which would seem not to influence A's and B's beliefs relating to p. So at the end of the dialogue we have the following situation with respect to the information p:

(2) a. A believes that p; B believes that p;
b. A believes that B believes that p; B believes that A believes that p.

In a shallow sense, p has become a shared belief: both participants have this belief and they both believe that the other has that belief. But studies of the logical foundations of communication tell us that participants in a dialogue should establish a *common ground* in a deeper sense. In their groundbreaking studies of common ground, Stalnaker and Lewis, among others, have suggested to define common ground in terms of *mutual beliefs*, explained as follows:

(3) p is a mutual belief of A and B iff:
- A and B believe that p;
- A and B believe that A and B believe that p;
- A and B believe that A and B believe that A and B believe that p;
and so on *ad infinitum*.

Clearly, the situation represented in (2) is a very poor approximation of this notion of common ground. Yet, intuitively, at the end of dialogue (1) the information that the next train to Amsterdam leaves at 9:17 seems to be *grounded*, i.e. to have been added to the common ground of A and B.

A technical problem presents itself here: the communicative acts expressed by the dialogue utterances create only finite iterations of belief of one dialogue participant about the beliefs of the other participant, as illustrated by (2); the full recursive nature of mutual beliefs cannot be achieved in this way in a dialogue of finite length.

In this paper we will describe a computational model of grounding where the establishment of common ground comes out as a consequence of successful communication, defined as the recognition of each other's intentions, plus two pragmatic principles, one concerning the way in which dialogue participants deal with expectations of being understood and believed; and one about the cumulative effects of feedback. The model, which does not require any specific grounding acts, is backed up by empirical observations from corpora of information-seeking and assistance dialogues.

This paper is organised as follows. Section 2 summarizes some existing views on grounding. Section 3 presents the conceptual model of grounding, based on dialogue analysis according to the framework of Dynamic Interpretation Theory (DIT, (Bunt, 2000)); section 4 presents our computational model of grounding, and Section 5 ends with concluding remarks.

2 Common Ground and Grounding

In Clark and Schaefer's model of grounding (Clark and Schaefer, 1989), participants in a dialogue try to establish for each utterance the mutual belief that the addressees have understood what the speaker meant. This is accomplished by the use of units called *contributions*. Contributions are divided into an acceptance and a presentation phase, so that every contribution, except for those that express negative evidence, has the role of accepting the previous contribution. A difficulty with this model is that its grounding criterion says that *"the contributor and the partners mutually believe that the partners have understood what the contributor meant"*. So the grounding *process* is conceived in terms of mutual beliefs. However, the central problem of grounding is precisely how mutual beliefs are established. Work based on this model includes its extension to human–computer interaction by Brennan and collaborators (Brennan, 1998; Cahn and S. E. Brennan, 1999), Li et al.'s model for multimodal grounding (Li et al., 2006), and Paek and Horvitz's formal theory of grounding (Paek and Horvitz, 2000).

In his influential computational model of grounding, Traum (1994) has introduced separate *grounding acts* which are used to provide communicative feedback and thereby create mutual beliefs. For this approach to work, Traum assumes that feedback acts are always correctly perceived and understood, therefore a dialogue participant does not need feedback about his feedback acts. This is an unwarranted assumption, however. Like any dialogue utterance, an utterance which expresses feedback can suffer from the addressee temporarily being disturbed by the phone, or by an aircraft flying over, or by noise on a communication channel; hence a speaker who performs a grounding act can never be sure that his act was performed successfully until he has received some form of feedback. A limitation and somewhat confusing aspect of this model is that it discusses the grounding of *utterances*, rather than the grounding of information conveyed by utterances through their semantic content.

(Matheson et al., 2000) use elements of Traum's model in their treatment of grounding from the Information State Update perspective. They represent grounded and ungrounded discourse units in the information state, and change their status from ungrounded to grounded through grounding acts. The dialogue act `Acknowledgement` is the only grounding act implemented; its main effect is to merge the information in the acknowledged discourse unit into the grounded information. They do not deal with cases of misunderstandings or cases where the user asks for acknowledgement. The model keeps only the last two utterances in the information state, so it is not clear what would happen if the utterance to be grounded is more than two utterances back – which we will argue to be the rule rather than the exception.

3 Grounding and Belief Strengthening

The addition of something to a common ground relies on evidence that the belief in question is mutually believed. The nature of such evidence depends on the communicative situation, for instance on whether the participants can see each other, and on whether they are talking about something they (both know that they) can both see. We restrict ourselves here to situations where grounding is achieved through verbal communication only, as in the case of telephone conversations, email chats, or spoken human-computer dialogue.

In the DIT framework, information can pass from one dialogue participant to another through mech-

anisms linked to understanding and believing each other. The first of these consists of the information state of the addressee of a dialogue act undergoing certain changes when he understands the corresponding dialogue behaviour. Understanding communicative behaviour is modeled as the addressee coming to believe that the preconditions hold which are characteristic for the dialogue acts that are expressed by that behaviour. For example, if A asks B a Yes/No-Question about a proposition p, then as a result of understanding this, B will know that A wants to know whether p, and that A thinks that B knows whether p. The second mechanism is that of belief adoption (a.k.a. 'belief transfer', Allen and Perrault, 1980). When A has asked B whether p, and B answers "Yes", then upon understanding this A will assume that B believes that p. In such a situation, A may be expected to believe B, so A also believes that p: he has *adopted* p.

To be sure that information is indeed transferred through the mechanisms of understanding and/or adoption, a speaker needs evidence of correct understanding of his communicative behaviour and of being believed. Feedback, positive or negative, provides information about an addressee's understanding and adoption of information.

Let us consider the transfer of information through understanding and adoption in some more detail, to see its contribution to grounding processes. In the following dialogue fragment, A initially contributes utterance du_1 which expresses an `Inform` act; let c_1 be the precondition that A *believes that* p, with p the propositional content of the act (the information that the next train is at 11:02). Successful communication should lead to c_1 as well as p at some point being in A's and B's common ground.

(4) du_1. A: *The next train is at 11:02.*
du_2. B: *At 11:02.*
du_3. A: *That's correct.*
du_4. B: *Okay thanks.*

How could A for example come to believe that p is mutually believed? First, he should have evidence that B understands his utterance du_1 and believes its content p. B's utterance du_2 can be taken to provide such evidence. So after du_2, A believes that B believes that p, and that B believes that A believes that p. However, A cannot be certain that B indeed believes that p, since in du_2 he also seems to offer that belief for confirmation. A's response du_3 gives that confirmation. At this point A does not yet know whether his utterance has reached B and was well understood. B's next contribution du_4 provides evidence for that; upon understanding du_4, A has accumulated the following beliefs:

(5) A believes that p
A believes that B believes that p
A believes that B believes that A believes that p
A believes that B believes that A believes that B believes that p
A believes that B believes that A believes that B believes that A believes that p

Although we see nested beliefs of some depth emerging, A is still a long way from believing that p is mutually believed – an infinitely long way, in fact. Clearly, continuing along this line could not lead to mutual beliefs in a finite amount of time. We therefore want to suggest a different explanation.

In natural face-to-face dialogue, speakers receive feedback while they are speaking as the participants give explicit and implicit feedback about their understanding of what is being said by means of facial expressions, head movements, direction of gaze, and verbal elements. In situations without visual contact, such as telephone dialogues or computer-mediated chatting, or in human-computer dialogue, a speaker often receives no feedback while speaking (or typing). This has the effect that, when a speaker has finished a turn, he does not know whether his contribution has been perceived, understood, and accepted. In a situation where "normal input-output" conditions hold (Searle, 1969), i.e. where participants speak the same language, have no hearing or speaking impairments, use communication channels without severe distortions, and so on, a speaker normally expects that the addressee perceives, understands and believes what is being said. We model this by the speaker having a doxastic attitude that we call *weak belief* that the addressee of his dialogue acts believes the preconditions and the content of the dialogue act to be true.[1] So after contributing an utterance that expresses a dialogue act with precondition c_1, the speaker A has the weak belief that B be-

[1] A weak belief is characteristically distinguished from a firm belief in that it is not inconsistent to weakly believe that p while at the same time having the goal to know whether p. In fact, the combination of such a goal and weak belief forms the preconditions of a CHECKQUESTION.

lieves that c_1. And similarly, in information-seeking dialogues, assistance dialogues, and other types of cooperative dialogue where the participants are expected to only provide correct information about the task at hand, if the utterance offers the information p about the task, then the speaker A also has the weak belief that B believes that c_1.

Of course, the assumptions of being understood and believed are not idiosyncratic for a particular speaker, but are commonly made by dialogue participants in cooperative dialogue in normal input-output conditions. B will therefore believe that A makes this assumption, so:

(6) B believes that A weakly believes that B believes that c_1. B believes that A weakly believes that B believes that p.

By the same token, A believes this to happen, hence:

(7) A believes that B believes that A weakly believes that B believes that c_1 and that p.

This line of reasoning can in principle be continued *ad infinitum*, leading to the conclusion that:

(8) Both A and B believe that it is mutually believed that A weakly believes that B believes that c_1 and that p.

In the example dialogue, this means in particular that, after contributing utterance du_1, A will among other things believe the following 'weak mutual beliefs' to have been established, 'weak' in the sense that the mutual belief contains a weak belief link:

(9) a. A believes that it is mutually believed that A weakly believes that B believes that c_1.

b. A believes that it is mutually believed that A weakly believes that B believes that p.

The first of these weak mutual beliefs comes from the expected understanding of du_1, the second from the expected adoption of the information that du_1 offered.

More generally, what we see happening with respect to grounding, is that for an agent to ground a belief, what he has to do is not so much extend a finite set of nested beliefs like (5) to an infinite set of nested beliefs of any depth, but to replace the weak belief link in believed mutual beliefs of the form

(10) *A believes that it is mutually believed that A* **weakly** *believes that B believes that q*

by an ordinary belief link, turning it into

(11) *A believes that it is mutually believed that A believes that B believes that q*

which is equivalent[2] to:

(12) *A believes that it is mutually believed that q*

So the question is **what evidence is necessary and sufficient to strengthen the weakest link** in certain 'weak mutual beliefs'.

We have suggested above that the evidence behind nested beliefs of the complexity of (5) is *necessary* but not sufficient. That it is indeed necessary can be seen from the following example.

(13) 1. A: Where should I insert the paper?
2. B: In the paper feeder.
3. A: The paper to be faxed.
4. B: What did you say?

This example illustrates the above remark that utterances which provide feedback on a previous utterance are themselves also in need of feedback in order to make sure that they contribute to the grounding process. With utterance 3, A explains what he meant by *the paper* in his previous utterance, thereby indicating that he's not sure that his question was correctly understood. In other words, utterance 2 apparently did not provide A with positive feedback relating to being understood. A would certainly not be allowed to ground, having insufficient evidence about the feedback that B has received up to this point in the dialogue. Hence at this point the process does not move into the direction of establishing a mutual belief about the preconditions of the question, let alone of the answer.

The issue of evidence being necessary and/or sufficient for strengthening the weakest link in a weak mutual belief is an empirical one. The case of (13) represents empirical evidence for the necessity of the evidence behind (5). Contrary to what we suggested above, empirical evidence in fact seems to

[2] This equivalence depends on the assumption that is known in epistemic logic as the Introspection axiom. According to this assumption, an agent believes his own beliefs, and in this case an agent also believes that he has a certain goal when he in fact has that goal. A precondition q of a dialogue act performed by a speaker A is always a property of A's state of beliefs and goals, hence *A believes that q* is equivalent to q. Moreover, all dialogue participants may be assumed to operate according to this assumption, hence *B believes that A believes that q* is equivalent to *B believes that q*.

show that the evidence of correct understanding that supports the beliefs represented in (5) is also *sufficient* for strengthening the weak mutual belief in (8). We express this observation as a pragmatic principle for the strengthening of the weakest link in a 'weak mutual belief'. The principle says that:

(14) a. A dialogue participant strengthens the weak belief link in a 'weak mutual mutual belief' concerning a precondition of a dialogue act that he has performed, when (1) he believes that the corresponding utterance was correctly understood; (2) he has evidence that: (2a) the other dialogue partner also believes that; and (2b) they both have evidence that they both have evidence that (1) and (2a) are the casse.

b. Like clause a., replacing "precondition of" by "task-related information, offered by", and replacing "correctly understood" by "believed".

We call (14) the *Strengthening Principle (SP)*. The SP may not seem very transparent at first; we will show its effect below, where we will see that it in fact comes down to a dialogue participant being able to ground preconditions or contents of a dialogue act when he has twice received positive feedback, namely positive feedback (possibly implicitly only) on the original utterance and positive feedback (again, possibly implicitly) on his response to that feedback act. In Morante (2007) and (Morante, forthcoming 2007) we provide ample empirical evidence for this principle, using corpora of both human-human and spoken human-computer dialogues; here we give just one example.

In dialogue (15), the SP predicts that B grounds the content of the first utterance when he successfully processes utterance 5 (second case of positive feedback). Indeed, it seems impossible for B to continue with utterance 6, expressing doubts about the grounded belief. By contrast, B could very well express such doubts in his previous turn, as (16) illustrates.

(15) 1. A: The next train is at 11:02.
2. B: At 11:02.
3. A: That's correct.
4. B: Okay thanks.
5. A: You're welcome.
6. B: *I thought it would be at 11:08.

(16) 1. A: The next train is at 11:02.
2. B: At 11:02.
3. A: That's correct.
4. B: I thought it would be at 11:08.

Since the only difference between (15) and (16) is the feedback that has been given by utterances 4 and 5, it must be the case that the evidence of correct understanding provided by these utterances makes the difference for grounding.

Limitations of space prevent us from going into the ways in which the various types of dialogue acts facilitate, speed up, or delay grounding in dialogue. See (Morante, forthcoming 2007) for a systematic discussion.

4 The DIT computational model of grounding

Our computational modeling of grounding, based on the strengthening of weak belief links in mutual beliefs, exploits the DIT structured context model and detailed analysis of feedback. The context model consists of several components, each representing a different type of information. The most relevant components to consider here are the *Linguistic Context*, the *Cognitive Context*, and the *Semantic Context*, which are defined as follows:

- Linguistic Context: a record of the dialogue up to this point, including verbatim representations of utterances as well as aspects of their syntactic, semantic and pragmatic analysis;
- Cognitive Context: information about the processing of utterances, notably about any problems in their interpretation or application;
- Semantic Context: information about the task, including nested beliefs about the dialogue partner's semantic context.

Evidence of correct understanding and of being believed, which triggers the application of the Strengthening Principle, is represented in the Cognitive Context. In order to see how the context updates, corresponding to understanding and believing each other, lead to the grounding of information, consider how the content of utterance 2 in the dialogue (17), *In the feeder*, is grounded.

(17) du_1. U: Where should I insert the paper?
du_2. S: In the feeder.
du_3. U: Should I put it in the bottom front tray?
du_4. S: No, in the open tray on top.
du_5. U: OK thanks .
du_6. S: You're welcome.
du_7. U: Goodbye.

We will represent the information that an utterance u was successfully processed (at all levels[3]) by agent Y as $Y^+(u)$, and the fact that agent X has evidence that agent Y succesfully processed that utterance as: $X : Y^+(u)$.[4]

Utterance du_3 in (17) shows a problem in understanding du_2 (represented by $U^-(du_2)$) in the form of a clarification question. As a result of recognizing this, S cancels the beliefs which reflected his expectation that du_2 would be understood without problems (the beliefs labeled ssc4 and ssc5 in Table 1).

Utterance du_5 provides evidence for U's understanding the answer du_4 as well as believing it, so successful processing of du_5 introduces the element $S : U^+(du_4)$ into S's cognitive context. Utterance du_6 likewise can be taken to provide evidence that the preceding utterance was well understood, so that leads to U's cognitive context containing the element $U : S^+(du_5)$. And similarly du_7 leads to S's cognitive context containing $S : U^+(du_6)$.

Due to the local nature that feedback usually has, especially positive feedback (and even more strongly *implicit* positive feedback), this process however does not build up the nested evidence of understanding and believing du_2 that we need for its content to be grounded via the Strenghtening Principle. The key to solving this problem can be found in the observation that, *when you get positive feedback on your last contribution to the dialogue, then that is evidence for you that the speaker thinks that you successfully processed his preceding contribution.*

For example, when you have been asked a question, then positive feedback on the answer that you give constitutes evidence that you had understood the question well. We call this phenomenon **Feedback Chaining**. It can be represented formally as:

(18) $S^+(du_i) \Rightarrow S : A^+(du_{i-1})$

(with S indicating Speaker and A Addressee). Negative feedback is of course a different story: understanding of a negative feedback act means for the addressee that he has to address the utterance that caused the negative feedback. In the example of (17) we see that S recognizes that du_3 signaled a problem with du_2 (item $S^+(du_3^{-2})$ in S's Cognitive Context).

Note that Feedback Chaining is something that all participants in a dialogue do and assume all participants to do. Utterance du_5 in the example dialogue therefore not only leads to the element $S : U^+(du_4)$ in S's cognitive context, saying that S has evidence that U successfully processed utterance du_4, but from applying Feedback Chaining to the new element in his cognitive context also to inferring that U has evidence that S successfully processed the utterance preceding du_4, hence that $S : U : S^+(du_3)$.

Table 1 shows some of the information in the linguistic context of the participant who has the speaker turn, and of the effects of what is said on the participants' cognitive and semantic contexts. Of the linguistic context it shows: (1) the verbatim form of each turn; (2) the speaker of that turn: (3) the chronological location of the turn; (4) the communicative functions of the dialogue acts performed in that turn, where for simplicity we only show the communicative functions that are relevant to the present discussion.

Feedback Chaining has the effect that dialogue acts that provide feedback, either explicitly or implicitly, have a non-local effect and allow dialogue participants to build up evidence about each other's evidence concerning the processing of utterances earlier in the dialogue, and at some stage this nested evidence meets the requirements of the Strengthening Principle. In the example dialogue, U can ground the preconditions of his question du_3 after utterance du_6 since he has evidence that du_3 was well understood (element ucc3 of his cognitive context), and that S has evidence that this is the case (el-

[3]DIT distinguishes several levels of feedback, namely those of paying attention, perception, understanding, evaluation, and application. The Feedback Chaining principle presented below is a simplification; in full it takes the various levels of feedback into account.

[4]Everywhere in this paper when we speak of 'feedback' we mean what in DIT is called *auto-feedback*, as opposed to *allo-feedback*. The former is concerned with information about the speaker's processing of dialogue utterances; the latter with the speaker's beliefs about the addressee's processing. For allo-feedback a similar chaining principle applies as the one described below for auto-feedback.

Table 1: Linguistic, Cognitive and Semantic contexts (slightly simplified) for dialogue (17)

LC = Linguistic Context; CC = Cognitive Context; SC = Semantic Context. c_{ki} stands for the preconditions of du_k; c_k for the semantic content of du_k. 'und' = understanding of previous utterance; 'exp' = expected; 'ad' = adoption; FC = Feedback Chaining; SP = Strengthening Principle. 'bel' = belief; 'wbel' = weak belief; 'mbel' - mutual belief.

	num	*source*	*S's context*	*num*	*source*	*U's context*
SC				usc1	prec	c_{1i}
LC				du_1	U	*Where should I insert the paper?* WH–QUESTION
CC	scc1	und	$S^+(du_1)$			
SC	ssc1 ssc2 ssc3	und exp und prec	bel(S, c_{1i}) bel(S, mbel(S, U, wbel(U, bel(S, c_{1i})))) bel(S, c_2)	usc2	exp und	bel(U, mbel(S, U, wbel(U, bel(S, c_{1i}))))
LC	du_2	S	*In the feeder.* WH–ANSWER(du_1)			
CC				ucc1	und	$U^-(du_2)$
SC	ssc4 ssc5	exp und exp ad	bel(S, mbel(S, U, wbel(S, bel(U, c_{2i})))) bel(S, mbel(S, U, wbel(S, bel(U, c_2))))	usc2 usc3	exp und exp ad	bel(U, mbel(S, U, wbel(S, bel(U, c_{2i})))) bel(U, mbel(S, U, wbel(S, bel(U, c_2))))
LC				du_3	U	*Should I put it in the bottom front tray?* NEG. FEEDBACK YN-QUESTION(du_2)
CC	scc2 scc3	und FC	$S^+(du_3^{-2})$ $S : U^-(du_2)$			
SC	 ssc6 ssc7 ssc8	 und exp und prec	cancellation of ssc4, ssc5 bel(S, c_{3i}) bel(S, mbel(S, U, wbel(U, bel(S, c_{3i})))) bel(S, c_4)	usc4	exp und	bel(U, mbel(S, U, wbel(U, bel(S, c_{3i}))))
LC	du_4	S	*No, in the open tray on top.* YN–ANSWER(du_3)			
CC				ucc2 ucc3	und FC	$U^+(du_4)$ $U : S^+(du_3)$
SC	 ssc9 ssc10	 exp und exp ad	 bel(S, mbel(S, U, wbel(S, bel(U, c_{4i})))) bel(S, mbel(S, U, wbel(S, bel(U, c_4))))	usc5 usc6 usc7	ad exp und exp ad	cancellation of usc2, usc3 bel(U, c_4) bel(U, mbel(S, U, wbel(S, bel(U, c_{4i})))) bel(U, mbel(S, U, wbel(S, bel(U, c_4))))
LC				du_5	U	*OK thanks.* POSITIVE FEEDBACK(du_4)
CC	scc4 scc5 scc6	und FC FC	$S^+(du_5)$ $S : U^+(du_4)$ $S : U : S^+(du_3)$			
LC	du_6	S	*You'e welcome* POSITIVE FEEDBACK(du_5)			
CC				ucc4 ucc5 ucc6 ucc7	und FC FC FC	$U^+(du_6)$ $U : S^+(du_5)$ $U : S : U^+(du_4)$ $U : S : U : S^+(du_3)$
SC				usc6	SP	bel(S, mbel(S, U, c_{3i}))
LC				du_7	U	*Goodbye* POSITIVE FEEDBACK(du_6)
CC	scc7 scc8 scc9 scc10 scc11	und FC FC FC FC	$S^+(du_7)$ $S : U^+(du_6)$ $S : U : S^+(du_5)$ $S : U : S : U^+(du_4)$ $S : U : S : U : S^+(du_3)$			
SC	ssc7	SP	bel(S, mbel(S, U, c_4))			

ement ucc7).[5] This is what we may call the *grounding of the utterance* by U.

From an intuitive point of view, S should perhaps also be able to ground utterance du_4. But does he in fact have evidence that U correctly understood that utterance? All that S has to go by is U's thanking and goodbye acts, taken to also signal that U believes to have understood S's answer du_4 successfully, but of course U may be wrong; U's belief cannot constitute solid evidence for S. If indeed we want utterances to be grounded in such situations, then we need an additional pragmatic principle saying that, when a dialogue participant expresses that he has successfully processed a dialogue utterance, then this will be believed unless there is evidence to the contrary. Since utterance du_7 provides no such counter-evidence, S may at this point indeed assume that U processed du_4 successfully.

Note that our model of grounding says that the content of du_4 is *not* grounded for U at the end of this dialogue. Doesn't that make it unsatisfactory for U to end the dialogue? We believe not: we have here an information-seeking dialogue, with U as the information seeking participant. As far as U is concerned, the dialogue may end as soon as he believes that his question (du_3), replacing his original question du_1 was well understood and has received an answer (du_4) that he believes. What more could an information-seeking agent want?

5 Concluding Remarks

We have presented a simple, empirically based computational model of grounding in dialogue as the result of the strengthening of weak mutual beliefs. These weak beliefs are created through the assumptions that participants in dialogue make about the understanding and acceptance of what they say when normal input-ouput conditions hold. A crucial role in this model is played by the Strengthening Principle, which says that a dialogue participant can strengthen a weak mutual belief when he has sufficient evidence about both participants' belief that the utterance, which caused the weak belief was understood and accepted by the other participant.

A proof of concept implementation of the grounding model, outlined here, has been integrated as part of the Dialogue Manager module in a speech-based information-extraction system (see (Keizer and Morante, 2007)). This implementation proves the technical validity of the grounding model, and forms a platform for experimenting for example with different forms of the Strenghtening Principle for different types of dialogue.

References

J. F. Allen and C. R. Perrault. 1980. Analyzing intention in dialogues. *Artificial Intelligence*, 15(3):143–178.

S. Brennan. 1998. The grounding problem in conversations with and through computers. In S.R. Fussell and R.J. Kreuz, editors, *Social and cognitive psychological approaches to interpersonal communication*, pages 201–225. Lawrence Erlbaum, Hillsdale, NJ.

H. Bunt. 2000. Dialogue pragmatics and context specification. In H. Bunt and W. Black, editors, *Abduction, Belief and Context in Dialogue*, pages 81–150. Benjamins, Amsterdam.

J. Cahn and S. Brennan. 1999. A psychological model of grounding and repair in dialog. In *Proc. AAAI FAll Symposium on Psychological Models of Communication in Collaborative Systems*, pages 25–33, North Falmouth, MA. AAAI.

H. Clark and E. Schaefer. 1989. Contributing to discourse. *Cognitive Science*, 13:259–294.

S. Keizer and R. Morante. 2007. Dialogue simulation and context dynamics for dialogue management. In *Proc. NODALIDA Conference*, Tartu, Estonia.

S. Li, B. Wrede, and G. Sagerer. 2006. A computational model of multi–modal grounding for human robot interaction. In *Proc. 7th SIGdial Workshop on Discourse and Dialogue*, pages 153–160. ACL, Sydney.

C. Matheson, M. Poesio, and D. R. Traum. 2000. Modelling grounding and discourse obligations using update rules. In *Proceedings NAACL*.

R. Morante. 2007. Computing meaning in interaction. PhD Thesis, Tilburg University.

T. Paek and Eric Horvitz. 2000. Toward a formal theory of grounding. Technical report MSR–TR–2000–40, Microsoft Research, Redmond, WA.

J. Searle. 1969. *Speech acts*. Cambridge University Press, Cambridge, UK.

D. Traum. 1994. *A Computational Theory of Grounding in Natural Language Conversation*. PhD Thesis. Dep. of Computer Science, University of Rochester.

[5] The other SP conditions are also satisfied since we assume that the attitude 'has evidence that', like the other doxastic attitudes that we use, is logically introspective (cf. footnote 2). Therefore $S^+(du_3) \Rightarrow S : S^+(du_3)$.