Conversational Acts and Non Sentential Utterances in Multilogue

Jonathan Ginzburg and Raquel Fernández
Department of Computer Science
King's College, London
The Strand, London WC2R 2LS
{ ginzburg,raquel }@dcs.kcl.ac.uk

1 Introduction

Dialogue—two person conversation—is by now a topic with an ever increasing theoretical, corpus-based, and implementational literature. In contrast, the study of *multilogue*—conversation with 3 or more participants—is still in its early stages. *The* fundamental issue in tackling multilogue is: to what extent do mechanisms motivated for dialogue (e.g. information states, protocols, update rules etc) scale up directly to multilogue?

In this paper, we present evidence relating to long distance (measured in turns) resolution of non sentential utterances (NSUs). The evidence shows significant asymmetries between dialogue and multilogue, as well as across NSU types:

- **Multilogue Short Answer (MSA) effect**: the lion's share of short answers *in dialogue* are adjacent to their antecedent; a substantial percentage of short answers *in multilogue* occur more than 5 turns away from their antecedent.
- **Non Short Answer Adjacency (NSAA) effect**: apart from short answers, all other types of NSU in dialogue *and* multilogue occur generally adjacent to their antecedent.

We will suggest that these effects gives interesting pointers to the existence of types of interaction which we dub *distributive interactions* where each participant actively participates and which scale up directly from dialogue (e.g. querying). These contrast with other types of interaction, dubbed here *collective interactions*, which do not, wherein one conversationalist can commit a group of others (e.g. grounding and assertion acceptance).

The paper is structured as follows: in section 2 we present data from the British National Corpus (BNC) concerning the resolution of NSUs in dialogue and multilogue. The main focus of this data is with the distance between antecedent and fragment.

In section 3, we survey recent work on multilogue and point to some issues it contends with that the present paper addresses.

In section 4 we develop a theoretical analysis of both the MSA effect and the NSA effect. Accounting for the MSA effect will have two components: (a) an essentially pragmatic explanation for the absence of long distance short answers and (b) a protocol for querying and assertion in multilogue that licenses long distance short answers. Explicating the NSA effect will involve postulating multilogue protocols for grounding and (assertion) acceptance, ones which do not directly mirror their dialogical counterparts. In this respect, they are consistent with the claims of (Garrod and Pickering forthcoming), that multilogue has more affinities with monologue than dialogue.

Finally, in section 5 we describe a pilot evaluation of three interaction protocols on BNC data, the essential issue being how well does (the context built up by) the protocol resolve short answers.

2 Long Distance NSUs in Dialogue and Multilogue: Data

This section describes the empirical evidence provided by corpus data extracted from the British National Corpus (BNC). The corpus we use in this investigation includes and extends the sub-corpus used in (Fernández and Ginzburg 2002). The current corpus is a sub-portion of the BNC conversational transcripts consisting of 14,315 sentences. It has been created by excerpting a 200-speaker-turn section from 54 BNC files. Of these files, 29 are transcripts of conversations between two dialogue participants, and 25 files are multilogue transcripts.

A total of 1285 NSUs were found in our sub-corpus. Table 1 shows the raw counts of NSUs found in the dialogue and multilogue transcripts, respectively.

	# NSUs	# BNC files
Dialogue	709	29
Multilogue	576	25
Total	1285	54

Table 1: Total of NSUs in Dialogue and Multilogue

All NSUs encountered within the corpus were manually classified according to the NSU typology presented in (Fernández and Ginzburg 2002). In order to be able to measure the distance between NSUs and their antecedents, all instances were additionally tagged with the sentence number of their antecedent utterance. Table 2 shows the distribution of NSU categories and their antecedent separation distance.

2.1 NSU-Antecedent Separation Distance

The last row in Table 2 shows the distribution of NSU-antecedent separation distances as percentages of the total of NSUs found. This allows us to see that about 87% of NSUs have a distance of 1 sentence (i.e. the antecedent was the immediately preceding sentence), and that the vast majority (about 96%) have a distance of 3 sentences or less.

Although the proportion of NSUs found in dialogue and multilogue is roughly the same (see Table 1 above), when taking into account the distance of NSUs from their antecedent, the proportion of long distance NSUs in multilogue increases radically: the longer the distance, the higher the proportion of NSUs that were found in multilogue. In fact, as Table 3 shows, NSUs that have a distance of 7 sentences or more appear exclusively in multilogue transcripts.

2.2 Short Answers

One striking result exhibited in Table 2 is the uneven distribution of long distance NSUs across categories. With a few exceptions, NSUs that have a distance of 3 sentences or more are exclusively short answers. Not only is the long distance phenomenon almost exclusively restricted to short answers, but the frequency of long distance short answers stands in strong

NSU Class	Example	Total	Distance 1	2	3	4	5	6	>6
Acknowledgment	*Mm mm.*	595	578	15	2				
Short Answer	*Ballet shoes.*	188	104	21	17	5	5	8	28
Affirmative Answer	*Yes.*	109	104	4			1		
Clarification Ellisis	*John?*	92	76	13	2	1			
Repeated Ack.	*His boss, right.*	86	81	2	3				
Rejection	*No.*	50	49	1					
Factual Modifier	*Brilliant!*	27	23	2	1	1			
Repeated Aff. Ans.	*Very far, yes.*	26	25	1					
Help Rejection	*No, my aunt.*	24	18	5		1			
Check Question	*Okay?*	22	15	7					
Filler	*... a cough.*	18	16	1		1			
Bare Mod. Phrase	*On the desk.*	16	11	4			1		
Sluice	*When?*	11	10	1					
Prop. Modifier	*Probably.*	11	10	1					
Conjunction Phrase	*Or a mirror.*	10	5	4	1				
		1285	1125	82	26	9	7	8	28
		%	87.6	6.3	2	0.6	0.5	0.6	2.1

Table 2: Total of NSUs sorted by Class and Distance

Distance	1	2	3	4	5	6	7	8	9	10	>10
Total	467	45	15	8	6	7	3	1	2	1	21
%	41	55	55	88	86	87	100	100	100	100	100

Table 3: Total and % of NSUs in Multilogue sorted by Distance

contrast to the other NSUs classes; indeed, over 44% of short answers have more than distance 1, and over 24% have distance 4 or more, like the last answer in the following example:

(1) Allan: How much do you think?
Cynthia: Three hundred pounds.
Sue: More.
Cynthia: A thousand pounds.
Allan: More.
Unknown: < *unclear* >
Allan: Eleven hundred quid apparently, just that. [BNC, G4X]

Another interesting result is that short answers do seem to be more common in multilogue than in dialogue—71% v. 29%. Also, the distance pattern exhibited by these two groups is strikingly different: Only 18% of short answers found in dialogue have a distance of more than 1 sentence, with all of them having a distance of at most 3, like the short answer in (2).

(2) Malcolm: ... cos what's three hundred and sixty divided by seven?
Anon 1: I don't know.
Malcolm: Yes I don't know either! < *laugh* >
Anon 1: Fifty four point fifty one point four. [BNC, KND]

In contrast, long distance short answers are common in multilogue: 56% of those found in multilogue have more than distance 1, while 37% have more than distance 3. This is

summarised in Table 4, which shows the total number of short answers found in dialogue and multilogue respectively, and the proportions sorted by distance over those totals.

Short Answers	Total #	1	2	3	> 3
Dialogue	54	82	9	9	0
Multilogue	134	44	11	8	37

Table 4: % over the totals found in dialogue and multilogue

It should be noted that all short answers in the corpus that have more than distance 6 a total of 28) appear in 4 BNC files (J8J, JK8, JYM, JJH). They are all transcripts of tutorials, training sessions or seminars, which exhibit a rather particular structure where several participants take turns to answer a question asked by the tutor or session leader.

(3) Anon1: How important is those three components and
what value would you put on them.
[...]
Anon3: Tone forty five. Body language thirty < *unclear* >.
Anon1: Thank you.
Anon4: Oh.
Anon1: Melanie.
Anon5: < *unclear* > twenty five.
Anon1: Yes.
Anon5: Tone of voice twenty five. [BNC, JYM]

3 Previous Work on Multilogue

3.1 Dialogue in Multiagent Systems

An area of research where one would expect to find proposals to deal with multiple party conversations is perhaps that of communication between autonomous software agents. However, even though many situations considered in multiagent systems do involve more than two agents, most interaction protocols are designed only for two participants at a time. This is the case of the protocol specifications provided by FIPA (Foundation for Intelligent Physical Agents) for agent communication language messages (FIPA 2003).

The FIPA interaction protocols (IP) are most typically designed for two participants, an initiator and a responder. Some IPs permit the broadcasting of a message to a group of addressees, and the reception of multiple responses by the original initiator. However, even though more than two agents participate in the communicative process, as (Dignum and Vreeswijk 2003) point out, such conversations can not be considered multilogue, but rather a number of parallel dialogues.

An illustrating example of such protocols is the Contract Net IP. In the FIPA Contract Net IP the initiator solicits proposals from other agents to perform a task by broadcasting a call for proposals to all other agents. Participants receiving the initiator's call generate responses, proposing to perform the task or, alternatively, refusing to propose. The initiator evaluates the received proposals and selects agents to perform the task; The selected agents will be sent an accept-proposal act and the remaining agents will receive a reject-proposal act.

It is not clear how actual multilogue conversations could be integrated in the framework provided by FIPA. However, some efforts are being made to investigate this and other issues in multiagent systems: (Dignum and Vreeswijk 2003), for instance, present a simple implementation intended as a first step towards empirically exploring the properties of multilogue conversation between autonomous software agents.

3.2 Natural Language Dialogue

Most studies of natural language dialogue, both theoretical and computational, focus on interaction between two conversation participants. Very little research has been done to investigate the requirements and idiosyncracies of conversation between multiple parties. One exception to this general trend is some of the recent work done by David Traum and colleagues, related to the Mission Rehearsal Exercise (MRE) Project (Traum and Rickel 2002).

The setting of the MRE project, used for training of Army personnel, is a virtual reality environment where multiple partners (including humans and other autonomous agents) engage in multi-conversation situations. The dialogue model used in the MRE project draws on the account proposed in (Poesio and Traum 1997) and its implementation in Information State-based systems (e.g. (Matheson, Poesio, and Traum 2000)). It uses a combination of both dialogue obligations and Questions Under Discussion, together with a model of the grounding process that involves recognition and construction of common ground units (CGUs) (see (Traum 2003)).

One of the fundamental issues, mentioned previously, which the MRE project has addressed is how the above elements of dialogue management change when considering multilogue situations. (Traum 2004) addresses parameters including turn management, conversation management, participant roles, initiative management. We focus here on a set of issues that revolve around obligations. Modelling of obligations and grounding becomes more complex when considering multilogue situations. The model of grounding implemented in the MRE project (inspired by (Poesio and Traum 1997)) can only be used in cases where there is a single initiator and responder. It is not clear what the model should be for multiple addresses: should the contents be considered grounded when any of the addressees has acknowledged them? Should evidence of understanding be required from every addressee?

Multilogue situations can also be problematic for models of obligations: for instance, when a question is asked, does every addressee have an individual obligation or, alternatively, there is an obligation assigned to the group, which is relieved when any addressee performs the required action? A related issue is the transferring of obligations: if an agent who has been asked a question redirects such question to another agent, does the first agent still have the obligation to answer? And if the second agent responds, is the obligation relieved? Traum suggests that the answer to some of these questions may depend on the type of activity performed (tutorials, ...).

We believe that the MSA and NSAA effects, described in section 2 offer some pointers concerning these problems.

4 Theoretical Discussion

4.1 Explanatory strategy for MSA Effect

The MSA effect has two parts:

(4) a. antecedent adjacency for short answers in dialogue.

b. Possibility or even probability of long distance short answers in multilogue.

We offer a 'pragmatic' account of (a), explain why this does not apply to multilogue, then in section 5 offer formal modelling for this possibility. In part, this will arise from the fact that querying in multilogue is *distributive*: once a question has been posed it can be addressed by more than a single CP

In dialogue certain commonly observed conditions will enforce adjacency between short answers and their interrogative antecedents.

(5) a. Questions have a simple, one phrase answer.

b. Questions can be answered immediately, without preparatory or subsequent discussion.

For multilogue (or at least certain genres thereof), (5) are less likely to be maintained: different CPs can supply different answers, even assuming that relative to each CP there is a simple, one phrase answer. The more CPs there are in a conversation, the smaller their common ground and the more likely the need for clarificatory interaction.

4.2 Explanatory strategy for NSAA Effect

In seeking to explain the NSAA effect some care is needed. One reason for this is that different classes of NSUs have wildly divergent frequencies. So while the generalization that Ack NSUs ('mm', 'OK' etc) satisfy antecedent adjacency seems amply supported (578/595 tokens satisfy this generalization), the analogous generalization for the class Sluice, for instance, seems far less well supported (only 11 tokens occured in this particular sampling.). In discussing the NSAA effect, then, we will restrict ourselves to a couple of types of relatively frequent types of NSU.[1]

[1]We should mention one systematic exception to the NSAA effect, even for dialogue. (Ginzburg 2003) discusses clarificatory 'why'-sluices *in dialogue*, which are argued to be different from all other clarificatory sluices. They do not typically involve an overt antecedent sub-utterance and whereas all other reprise sluices invariably occur adjacently to the utterance of which they seek clarification, 'why'-sluices naturally occur several turns away:

(6) (i) Norrine: When is the barbecue, the twentieth? (pause) Something of June
Chris: Thirtieth.
Norrine: A Sunday.
Chris: Sunday.
Norrine: Mm.
Chris: Why? (= *Why do you ask when the barbecue is*) (BNC)

(Ginzburg 2003) suggests that *who, when, where* and indeed virtually all other *wh*-phrases used in clarification requests query a constituent of what one might call the nucleus of a given content (bold face in (ii,iii), whereas Why_{meta} outscopes this nucleus to query at the illocutionary level (bold face in (iv,v)):

We limit discussion to the classes Ack (acknowledgements such as 'mm', 'OK') and CE (clarification ellipsis e.g. 'Bo?' 'Who?', 'Why?' used in clarification requests.).

The basic conclusion we draw from the (essentially) obligatory adjacency of acknowledgement and CE NSUs to their antecedents is that these classes of dialogue moves in multilogue reflect what we will call *collective* dialogue moves—one response by a CP is sufficient to ground an utterance for an entire audience. Conversely, if a clarification request is not made by the CP whose turn succeeds a given utterance, then on the whole subsequent CPs should not pose a CR. AN example of such a protocol is sketched in section 5.

5 Protocols for Dialogue and Multilogue

In this section we sketch some conversational interaction protocols for multilogue. These serve to underpin our account of NSU resolution. We start by discussing the latter and the existing work on dialogue which constitutes our starting point.

5.1 Dialogue Interaction

Information States We assume information states of the kind developed in the KOS framework and implemented in the GODIS/IBIS systems (see e.g. (Ginzburg 1996; Larsson 2002; Ginzburg forthcoming)). On this view each dialogue participant's view of the common ground is structured by a number of attributes including the following three: FACTS: a set of facts representing the shared assumptions of the CPs, LatestMove: the most recent grounded move, and QUD ('questions under discussion'): a set—often taken to be structured as a stack—consisting of the currently discussable questions.

Querying and Assertion Both querying and assertion involve a question becoming maximal in the querier/asserter's QUD: the posed question q for a query where q is posed, the polar question $\lambda\{\ \}p$ for an assertion where p is asserted. Roughly, the responder can subsequently either choose to start a discussion (of q or p?) or, in the case of assertion, to update her FACTS structure with p. A dialogue participant can downdate $q/p?$ from QUD when, as far as her (not necessarily public) goals dictate, sufficient information has been accumulated in FACTS. The querying/assertion protocols (in their most basic form) are summarised as follows:

(7) cooperative query exchange

1. `LatestMove.Cont = Ask(A,q):  IllocProp`
2. `A: push q onto QUD; release turn`
3. `B: push q onto QUD; take turn; make q-specific utterance`[2]

(ii) A: Did Bo leave? B: Who?/When?/Where?

(iii) $\lambda b/t/l Ask(A, ?\mathbf{Leave}(\mathbf{b}, \mathbf{t}, \mathbf{l})$

(iv) A: Did Bo leave? B: Why?

(v) $\lambda r.Cause(r, \mathbf{Ask}(\mathbf{A}, ?Leave(b, t, l)))$

[2] An utterance whose content is either a proposition p About q/or a question q_1 on which q Depends

(8) cooperative assertion exchange

1. LatestMove.Cont = Assert(A,p): IllocProp
2. A: push p? onto QUD, release turn
3. B: push p? onto QUD, take turn; ⟨ Option 1: Discuss p?, Option 2: Accept p ⟩

(9)
1. LatestMove.Cont = Accept(B,p) : IllocProp
2. B: increment FACTS with p; pop p? from QUD;
3. A: increment FACTS with p; pop p? from QUD;

NSU Resolution We assume the account of NSUs developed in (Ginzburg and Sag 2001). The essential idea they develop is that NSUs get their main predicates from context, specifically via unification with the question that is currently *under discussion*, an entity dubbed the *maximal question under discussion* (MAX-QUD). NSU resolution is, consequently, tied to conversational topic, viz. the MAX-QUD.[3]

5.2 Querying and Assertion in Multilogue

How to extend protocols such as (7) and (8) to multilogue? There are obviously a number of options one can take and indeed it is not implausible to assume that different genres make use of distinct options.

We view the MSA effect as providing some evidence towards the distributed nature of querying in multilogue. In other words, a posed question is available to each of the participants of the conversation, after it has been introduced for discussion. To effect this one needs to modify somewhat the assumption that QUD is a stack. To see this, consider the constructed example in (10). The question posed by A can be addressed by B, C, and D using a short answer, the existence of other contributions notwithstanding. On the other hand, a propositional lexeme such as 'no' does seem to relate to the latest contribution,[4] as illustrated in the contrast between C.ii and D.ii:

(10) A: Who should we invite for the conference?

B: Svetlanov.

C: (i) Zhdanov / (ii) No (=Not Svetlanov), Zhdanov

D: (i) Gergev/ (ii) No (= Not Zhdanov, ≠ Not Svetlanov), Gergev / Not Svetlanov, Gergev.

In order to achieve a distributive effect, the ordering on QUD is modified as follows: questions introduced into QUD by queries do not get subsumed by their answers; questions introduced into QUD by assertions stack straightforwardly among each other. This means that the multilogue protocol for querying is a direct generalization of the dialogue case:

[3] The resolution of NSUs, on this approach, involves one other parameter, an antecedent sub-utterance dubbed the *salient-utterance* (SAL-UTT). This plays a role similar to the role played by the *parallel element* in higher order unification–based approaches to ellipsis resolution (see e.g. (Pulman 1997). For current purposes, we limit attention to the MAX-QUD as the nucleus of NSU resolution.

[4] modulo insertion sequences

(11) Querying for a conversation involving { A,B,C_1,...,C_n }

1. `LatestMove.Cont = Ask(A,q): IllocProp`
2. `A: push q onto QUD; release turn`
3. C_i: `push q onto QUD`
4. `B: push q onto QUD; take turn; make q-specific utterance`

The protocol for assertion requires two modifications: one simply reflects the modified ordering on QUD. The other involves the communal nature of the acceptance option. Acceptance by one CP can count as acceptance by all other addressees of an assertion. There is an obvious rational motivation for this, given the difficulty of a CP constantly monitoring an entire audience (when this consists of more than one addressee) for acceptance signals. It also enforces quick reaction to an assertion—anyone wishing to dissent from p must get their reaction in early i.e. immediately following the assertion since further discussion of p? is not countenanced if acceptance takes place. The latter can happen of course as a consequence of a dissenter not being quick on their feet; such cases are analyzed here as two CPs competing for the post-assertoric turn.

(12) Querying for a conversation involving {A,B,C_1,...,C_n}

1. `LatestMove.Cont = Assert(A,p): IllocProp`
2. `A: p? becomes assertion--issue maximal in QUD`[5]`, release turn`
3. C_i: `p? becomes assertion--issue maximal in QUD`
4. `B: p? becomes assertion--issue maximal in QUD, take turn;` ⟨ `Option 1: Discuss p?, Option 2: Accept p` ⟩

(13)
1. `LatestMove.Cont = Accept(B,p) : IllocProp`
2. `B: increment FACTS with p; pop p? from QUD;`
3. C_i: `increment FACTS with p; pop p? from QUD;`
4. `A: increment FACTS with p; pop p? from QUD;`

We believe acknowledgement in grounding interaction works in a similar fashion.

6 Evaluation

In this section we consider three possible protocols and evaluate them on a portion of the BNC.

- Baseline Protocol: it keeps only the last contribution made.
- Stack Protocol: all contributions are kept as a LIFO stack.
- Partial Order (refined stack) Protocol: questions introduced by queries do not get subsumed by their answers; questions introduced by assertions stack straightforwardly.

[5] That is, p? subsumes all q in QUD introduced by previous assertions.

We evaluated the protocols on a small sub-corpus, created by extracting a 50-speaker-turn section from 10 BNC files (3 transcripts of free conversation, 1 interview, 1 discussion, 1 business meeting, 1 seminar, and 3 trainning sessions/lectures). We found a total of 41 short answers. From these, 22 were answers to an immediately preceeding question. All short answers but 1 could be handled by either a stack or the refined stack protocol. The example that could not be handled by any protocol was a dialogue where there is a clarification question which is not answered:

(14) Helen: Who did Eve have on her plate?
Clare: Plate?
Mrs Tiggywinkle. [BNC, KCD]

A possible example that might provide evidence for the need of the refined stack protocol is (15), where the utterances in italics are uttered simultaneously, indicating that both the assertion (the last answer or a previous assertion?) and the initial question are available.

(15) Joyce: You had what?
He had some stuff nicked, a ski jacket which cost me seventy five quid it were half, the rest it should of been a hundred and fifty.
Ann: Nicked?
Joyce: Nicked
Alec: Mm
Joyce: Pinched
Ann: Aargh
Joyce: Pair of football boots
Ann: *nicked, ooh*
Joyce: *And a pair of trainers.* [BNC, KB2]

Protocols	Success rate
Baseline	53.6
Stack	97.5
Refined Stack	97.5

Table 5: Evaluation Results

7 Conclusions and Future Work

This paper attempts to provide empirical evidence for the fundamental question research on multilogue—conversation involving three or more persons—must engage in: what commonalities and differences exist between dialogue and multilogue?

In this paper we have discussed two empirical findings that concern the resolution of NSUs in dialogue and multilogue. The MSA effect indicates that whereas antecedents for short answers are almost invariably adjacent turnwise, a substantial percentage of antecedents for short answers in multilogue are found several turns away. We have argued that the MSA

effect suggests that querying (in some genres of multilogue at least) is a distributive across the CPs. We suggested that this can be captured in a QUD-based framework by a simple modification of the structure of QUD.

The NSAA effect contrasts the existence of long distance antecedents for short answers with the essential absence of such antecedents for other classes of NSUs, in particular acknowledgements and clarification requests. This suggests the existence of collective acts in which an act by a single CP conveys common ground status on an utterance/assertion for an entire audience.

We have also provided results of preliminary evaluation for several multilogue protocols relative to their ability to provide antecedents for short answer resolution. We intend to carry more extensive evaluation along these lines of these and other protocols we will develop.

References

Clark, H. (1996). *Using Language*. Cambridge: Cambridge University Press.

Dignum, F. P. and G. A. Vreeswijk (2003). Towards a testbed for multi-party dialogues. In *Proceedings of the first International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS 2003)*.

Fernández, R. and J. Ginzburg (2002). Non-sentential utterances: A corpus study. *Traitement automatique des languages. Dialogue 43*(2), 13–42.

FIPA (2003). The foundation for intelligent physical agents. interaction protocol specifications. http://www.fipa.org.

Garrod, S. and M. Pickering (forthcoming). Toward a mechanistic psychology of dialogue. *Behavioural and Brain Sciences*. Forthcoming.

Ginzburg, J. (1996). 1996. Interrogatives: Questions, facts, and dialogue. In Shalom Lappin, editor, *Handbook of Contemporary Semantic Theory*. Blackwell, Oxford.

Ginzburg, J. (2003). Disentangling public from private meaning. In J. van Kuppevelt and R. Smith (Eds.), *Current and New Directions in Discourse and Dialogue*, pp. 183–212. Kluwer.

Ginzburg, J. (forthcoming). *Semantics and Interaction in Dialogue*. CSLI Publications and University of Chicago Press, Stanford: California. Draft chapters available from http://www.dcs.kcl.ac.uk/staff/ginzburg.

Ginzburg, J and I. Sag (2001). *Interrogative Investigations*. CSLI Publications, Stanford, California.

Larsson, S. (2002). *Issue based Dialogue Management*. Ph. D. thesis, Gothenburg University.

Matheson, C., M. Poesio, and D. Traum (2000, April). Modeling grounding and discourse obligations using update rules. In *Proceedings of the 1st Annual Meeting of the North American Chapter of the ACL*, Seattle.

Poesio, M. and D. Traum (1997). Conversational actions and discourse situations. *Computational Intelligence 13*(3).

Pulman, S. (1997). Focus and Higher Order Unification. *Linguistics and Philosophy*, 20.

Traum, D. (2003). Semantics and pragmatics of questions and answers for dialogue agents. In H. Bunt (Ed.), *Proceedings of the 5th International Workshop on Computational Semantics*, Tilburg, pp. 380–394. ITK, Tilburg University.

Traum, D. (2004). Issues in multi-party dialogues. In F. Dignum (Ed.), *Advances in Agent Communication*, pp. 201–211. Springer-Verlag.

Traum, D. and J. Rickel (2002). Embodied agents for multi-party dialogue in immersive virtual world. In *Proceedings of the first International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS 2002)*, pp. 766–773.