Contextual phenomena and thematic relations in database QA dialogues: results from a Wizard-of-Oz Experiment

Núria Bertomeu, Hans Uszkoreit

Saarland University

Saarbrücken, Germany

uszkoreit|bertomeu@coli.uni-sb.de

Anette Frank, Hans-Ulrich Krieger and Brigitte Jörg

German Research Center of Artificial Intelligence

Saarbrücken, Germany

frank|krieger|joerg@dfki.de

Abstract

Considering data obtained from a corpus of database QA dialogues, we address the nature of the discourse structure needed to resolve the several kinds of contextual phenomena found in our corpus. We look at the thematic relations holding between questions and the preceding context and discuss to which extent thematic relatedness plays a role in discourse structure.

1 Introduction

As pointed out by several authors (Kato et al., 2004), (Chai and Ron, 2004), the information needs of users interacting with QA systems often go beyond a single stand-alone question. Often users want to research about a particular topic or event or solve a specific task. In such interactions we can expect that the individual user questions will be thematically connected, giving the users the possibility of reusing part of the context when formulating new questions.

That users implicitly refer to and even omit material which can be recovered from the context has already been replicated in several Wizard-of-Oz experiments simulating natural language interfaces to databases, (Carbonell, 1983), (Dahlbäck and Jönsson, 1989), the most frequent contextual phenomena being ellipsis, anaphora and definite descriptions.

A big challenge for interactive QA systems is, thus, the resolution of contextual phenomena. In order to be able to do so a system has to keep track of

the user's focus of attention as the interaction proceeds. The attentional state at a given point in the interaction is given by the discourse structure. An open issue, however, is the nature of the discourse structure model needed in a QA system. Ahrenberg et al. (1990) argue that the discourse structure in NL interfaces is, given the limited set of actions to be performed by the system and the user, simpler than the one underlying human-human dialogue. Upon Ahrenberg et al. (1990) this is given by the discourse goals, rather than the overall goals of the user, as is the case in task-oriented dialogues, (Grosz and Sidner, 1986). Following Ahrenberg et al. (1990), the QA discourse is structured in segments composed by a pair of initiative-response units, like question-answer, or question-assertion, in the absence of an answer. A segment can be embedded in another segment if it is composed by a clarification request and its corresponding answer. The local context of a segment is given by the immediately preceding segment. Upon Ahrenberg et al. (1990), the latter reliably limits up the search space for antecedents of anaphoric devices and ellipsis. However, as we will see, there are few cases where the antecedents of contextual phenomena are to be found beyond the immediately preceding segments. This suggests that a more complex approach to discourse structure for QA systems is needed.

In more recent studies of interactive QA special attention has been paid to the thematic relatedness of questions, (Chai and Ron, 2004), (Kato et al., 2004). Chai and Ron (2004) propose a discourse modeling for QA interactions in which they keep track of thematic transitions between questions. Although

the applications of tracking thematic transitions between questions have not been investigated in depth, Sun and Chai (2006) report on an experiment which shows that the use of a model of topic transitions based on Centering Theory improves query expansion for context questions. However, these previous studies on the thematic relations between questions are not based on collections of interactive data, but on questions centered around a topic that were collected in non-interactive environments. This means that they do not consider the answers to the questions, to which following questions can be related.

This paper presents data on different kinds of contextual phenomena found in a corpus of written natural language QA exchanges between human users and a human agent representing an interactive information service. We address two issues: the kinds and frequencies of thematic relations holding between the user questions and the preceding context, on the one hand, and the location of antecedents for the different contextual phenomena, on the other. We also discuss the question whether thematic relations can contribute to determine discourse structure and, thus, to the resolution of the contextual phenomena.

In the next section we present our data collection and the aspects of the annotation scheme which are relevant to the current work. In section 3 we present data regarding the overall thematic cohesion of the QA sessions. In section 4 we report on data regarding the co-occurrence of discourse phenomena and thematic relations and the distance between the phenomena and their antecedents. Finally, we discuss our findings with regard to their relevance with respect to the nature of discourse structure.

2 Corpus and methodology

2.1 Experimental set-up

In order to obtain a corpus of natural QA interactions, we designed a Wizard-of-Oz experiment. The experiment was set up in such a way that the exchanges between users and information system would be as representative as possible for the interaction between users and QA systems. We chose an ontology database instead of a text based closed domain QA system, however, because in order to simulate a real system short time responses were needed.

30 subjects took part in the experiment, which consisted in solving a task by querying LT-WORLD, an ontology containing information about language technology¹, in English. The modality of interaction was typing through a chat-like interface.

Three different tasks were designed: two of them concentrated on information browsing, the other one on information gathering. In the first task subjects had to find three traineeships at three different projects in three different institutions each on a different topic, and obtain some information about the chosen projects, like a contact address, a description, etc. In the second task, subjects had to find three conferences in the winter term and three conferences in the summer term, each one on a different topic and they had to obtain some information on the chosen conferences such as deadline, place, date, etc. Finally, the third task consisted of finding information for writing a report on European language technology in the last ten years. To this end, subjects had to obtain quantitative information on patents, organizations, conferences, etc.

The Wizard was limited to very few types of responses. The main response was answering a question. In addition, she would provide intermediate information about the state of processing if the retrieval took too long. She could also make statements about the contents of the database when it did not contain the information asked for or when the user appeared confused about the structure of the domain. Finally, she could ask for clarification or more specificity when the question could not be understood. Yet the Wizard was not allowed to take the initiative by offering information that was not explicitly asked for. Thus all actions of the Wizard were directly dependent on those of the user.

As a result we obtained a corpus of 33 logs (30 plus 3 pilot experiments) containing 125.534 words in 2.534 turns, 1.174 of which are user turns.

2.2 Annotation scheme

The corpus received a multi-layer annotation² consisting of five levels. The levels of turns and part-of-speech were automatically annotated. The level of turns records information about the speaker and time

¹See <http://www.lt-world.org>.

²We employed the annotation tool MMAX2 developed at EML Research, Heidelberg.

stamp. For the other levels - the questions level, the utterances level, and the entities level - a specific annotation scheme was developed. For these, we only explain the aspects relevant for the present study.

2.2.1 Questions

This level was conceived to keep track of the questions asked by the user which correspond to queries to the database. With the aim of annotating thematic relatedness between questions we distinguished two main kinds of thematic relations: those holding between a question and a previous question, *quest(ion)-to-quest(ion)-rel(ation)*, and those holding between a question and a previous answer, *quest(ion)-to-answ(er)-rel(ation)*.

Quest-to-quest-rels can be of the following types:

- *refinement* if the current question asks for the same type of entity as some previous question, but the restricting conditions are different, asking, thus, for a subset, superset or disjoint set of the same class.

(1) US: How many projects on language technologies are there right now?

US: How many have been done in the past?

- *theme-entity* if the current question is about the same entity as some previous question.

(2) US: Where will the conference take place?

US: What is the dead-line for applicants?

- *theme-property* if the current question asks for the same property as the immediately preceding question but for another entity.

(3) US: Dates of TALK project?

US: Dates of DEREKO?

- *paraphrase* if the question is the rephrasing of some previous question.

- *overlap* if the content of a question is subsumed by the content of some previous question.

We distinguish the following *quest-to-answ-rels*:

- *refinement* if the current question asks for a subset of the entities given in the previous answer.

(4) LT: 3810.

US: How many of them do research on language technology?

- *theme* if the current question asks about an entity first introduced in some previous answer.

(5) LT: Semaduct, ...

US: What language technology topics does the Semaduct project investigate?

Although Chai and Jin (2004) only consider transitions among questions in dialogues about events, most of our relations have a correspondence with theirs. *Refinement* corresponds to their *constraint refinement*, *theme-property* to their *participant-shift*, and *theme-entity* to their *topic exploration*.

2.2.2 Utterances

Utterances are classified according to their speech-act: *question*, *answer*, *assertion*, or *request*. Our annotation of discourse structure is identical in spirit to the one proposed by Ahrenberg et al. (1990). A segment is opened with a user question to the database and is closed with its corresponding answer or an assertion by the system. Clarification requests and their corresponding answers form segments which are embedded in other segments. Requests to wait and assertions about the processing of a question are also embedded in the segment opened by the question.

Fragmentary utterances are annotated at this level. We distinguish between fragments with a full linguistic source, fragments with a partial source, and fragments showing a certain analogy with the source. The first group corresponds to fragments which are structurally identical to the source and can, thus, be resolved by substitution or extension.

(6) US: Are there any projects on spell checking in Europe in the year 2006?

US: And in the year 2005?

Fragments with a partial source implicitly refer to some entity previously introduced, but some inference must be done in order to resolve them.

(7) US: How is the contact for that project?

US: Homepage?

The last group is formed by fragments which show some kind of parallelism with the source but which cannot be resolved by substitution.

- (8) US: Which conferences are offered in this winter term in the subject of English language?
 US: Any conferences concerning linguistics in general?

2.2.3 Reference

We distinguish the following types of reference to entities: identity or co-reference, subset/superset and bridging.

Co-reference occurs when two or more expressions denote the same entity. Within this group we found the following types of implicit co-referring expressions which involve different degrees of explicitness: elided NPs, anaphoric and deictic pronouns, deictic NPs, and co-referent definite NPs. Elided NPs are optional arguments, that is, they don't need to be in the surface-form of the sentence, but are present in the semantic interpretation. In (9) there is an anaphoric pronoun and an elided NP both referring to the conference *Speech TEK West 2006*.

- (9) US: *The Speech TEK West 2006*, when does it take place?
 LT: 2006-03-30 - 2006-04-01.
 US: Until when can I hand in a paper []?

Bridging is a definite description which refers to an entity related to some entity in the focus of attention. The resolution of bridging requires some inference to be done in order to establish the connection between the two entities. In example (2) in subsection 2.2.1 there is an occurrence of bridging, where *the dead-line* is meant to be the dead-line of the conference currently under discussion.

Finally, subset/superset reference takes place when a linguistic expression denotes a subset or superset of the set of entities denoted by some previous linguistic expression. Subset/superset reference is sometimes expressed through two interesting contextual phenomena: nominal ellipsis³, also called semantic ellipsis, and one-NPs⁴. Nominal ellipsis occurs within an NP and it is namely the noun what

³Note, however, that nominal ellipsis does not necessarily always denote a subset, but sometimes it can denote a disjoint set, or just lexical material which is omitted.

⁴One-NPs are a very rare in our corpus, so we are not considering them in the present study.

is missing and must be recovered from the context. Here follows an example:

- (10) US: Show me *the three most important*.

3 Thematic follow-up

When looking at the thematic relatedness of the questions it's striking how well structured the interactions are regarding thematic relatedness. From 1047 queries to the database, 948 (90.54%) follow-up on some previous question or answer, or both. Only 99 questions (9.46%) open a new topic. 725 questions (69.25% of the total, 76.48% of the connected questions) are related to other questions, 332 (31.71% of the total, 35.02% of the connected questions) are related to answers, and 109 (10.41% of the total, 11.49% of the connected questions) are connected to both questions and answers. These numbers don't say much about how well structured the discourse is, since the questions could be far away from the questions or answers they are related to. However, this is very seldom the case. In 60% of the cases where the questions are thematically connected, they immediately follow the question they are related to, that is, the two questions are consecutive⁵. In 16.56% of the cases the questions immediately follow the answer they are related to. 74.58% of the questions, thus, immediately follow up the question or/and answer they are thematically related to⁶.

Table 1 shows the distribution of occurrences and distances in segments for each of the relations described in subsection 2.2.1. We found that the most frequent question-to-question relation is *theme-entity*, followed by the question-to-answer relation *theme*. As you can see, for all the relations except *theme*, most occurrences are between very close standing questions or questions and answers, most of them holding between consecutive questions or questions and answers. The occurrences of the relation *theme*, however, are distributed along a wide range of distances, 29.70% holding between questions and answers that are 2 and 14 turns away from

⁵By consecutive we mean that there is no intervening query to the database between the two questions. This doesn't imply that there aren't several intervening utterances and turns.

⁶9 questions are consecutive to the question and answer they are related to, respectively, that's why the total percentage of related consecutive questions is not 76.56%.

	REF. Q.	THEME E. Q.	THEME P. Q.	PARA. Q.	OVERL. Q.	REF. A.	THEME A.
TOTAL	74 (7.80%)	338 (35.65%)	107 (11.29%)	174 (18.35%)	29 (3.06%)	29 (3.06%)	303 (31.96%)
1 SEGM.	88.73%	81.65%	100%	60.92%	78.57%	83.34%	46.39%
2 SEGM.	5.63%	1.86%	0%	8.09%	21.43%	13.33%	10.20%

Table 1: Occurrences of the different thematic relations

REL. / PHEN.	THEME E. Q.	THEME P. Q.	THEME A.	REF. Q.	REF. A.	CONNECTED	TOTAL
FRAGMENT	53 (54.08%)	17 (16.32%)	3 (3.06%)	21 (21.42%)	0	97 (85.08%)	114
BRIDGING	40 (74.07%)	0	3 (5.55%)	1 (1.85%)	0	54 (58.69%)	92
DEFINITE NP	26 (78.78%)	0	4 (12.21%)	2 (6.10%)	0	33 (66%)	50
DEICTIC NP	19 (51.35%)	0	13 (35.13%)	2 (5.40%)	1 (2.70%)	37 (78.72%)	47
ANAPHORIC PRON.	13 (39.39%)	2 (6.06%)	10 (30.30%)	0	5 (15.15%)	33 (39.75%)	83
DEICTIC PRON.	2 (75%)	0	1 (25%)	0	0	3 (25%)	12
ELIDED NP	9 (69.23%)	0	2 (15.38%)	0	0	13 (61.90%)	21
NOMINAL ELLIPSIS	0	1 (7.69%)	6 (46.15%)	1 (7.69%)	5 (38.46%)	13 (81.25%)	16

Table 2: Contextual phenomena and the thematic relations holding between the questions containing them and the questions or answers containing the antecedents.

each other. This is because often several entities are retrieved with a single query and addressed later on separately, obtaining all the information needed about each of them before turning to the next one. We found also quite long distances for paraphrases, which means that the user probably forgot that he had asked that question, since he could have also scrolled back.

These particular distributions of thematic relations seem to be dependent on the nature of the tasks. We found some differences across tasks: the information gathering task elicited more refinement, while the information browsing tasks gave rise to more theme relations. It is possible that in an interaction around an event or topic we may find additional kinds of thematic relations and different distributions. We also observed different strategies among the subjects. The most common was to ask everything about an entity before turning to the next one, but some subjects preferred to ask about the value of a property for all the entities under discussion before turning to the next property.

4 Contextual phenomena: distances and thematic relatedness

There are 1113 user utterances in our corpus, 409 of which exhibit some kind of discourse phenomenon, i.e., they are context-dependent in some way. This amounts to 36.16% of the user utterances, a pro-

portion which is in the middle of those found in the several corpora analyzed by Dahlbäck and Jönsson (1989)⁷. The amount of context-dependent user utterances, as Dahlbäck and Jönsson (1989) already pointed out, as well as the distribution of the different relations among questions and answers explained above, may be dependent on the nature of the task attempted in the dialogue.

Table 2 shows the distribution of the most frequent thematic relations holding between the questions containing the contextual phenomena considered in our study and the questions or answers containing their antecedents. The rightmost column shows the number of occurrences of each of the contextual phenomena described in subsection 2.2.3. The second column on the right shows the number of occurrences in which the antecedent is located in some previous segment and the question containing the contextual phenomenon is related through a thematic relation to the question or answer containing the antecedent. The percentages shown for each phenomenon are out of the total number of its occurrences. The remaining columns show frequen-

⁷They found a high variance according to the kind of task carried out in the different dialogues. Dialogues from tasks where there was the possibility to order something contained a higher number of context-dependent user initiatives, up to 54.62%, while information browsing dialogues contained a smaller number of context-dependent user initiatives, 16.95% being the lowest amount found.

cies of co-occurrence for each of the phenomena and thematic relations. The percentages shown for each phenomenon are out of the total number of its connected occurrences.

For the majority of investigated phenomena we observe that most questions exhibiting them stand in a thematic relation to the question or answer containing the antecedent. Although there may be several intermediate turns, the related questions are almost always consecutive, that is, the segment containing the contextual phenomenon immediately follows the segment containing the antecedent. In the remainder of the cases, the contextual phenomenon and its antecedent are usually in the same segment.

However, this is not the case for deictic and anaphoric pronouns. In most cases their antecedents are in the same segment and even in the same utterance or just one utterance away. This suggests that pronouns are produced in a more local context than other phenomena and their antecedents are first to be looked for in the current segment.

For almost all the phenomena the most frequent relation holding between the question containing them and the question or answer containing the antecedent is the question-to-question relation *theme-entity*, followed by the question-to-answer relation *theme*. This is not surprising, since we refer back to entities because we keep speaking about them.

However, fragments and nominal ellipsis show a different distribution. Fragments are related to their sources through the question-to-question relations *theme-property* and *refinement*, as well. Regarding the distribution of relations across the three different types of fragments we distinguish in our study, we find that the relations *refinement* and *theme-property* only hold between fragments with a full source and fragments of type analogy, and their respective sources. On the other hand, practically all fragments with a partial-source stand in a *theme-entity* relation to their source. Questions containing nominal ellipsis are mostly related to the preceding answer both through the relations *theme* and *refinement*.

4.1 Antecedents beyond the boundaries of the immediately preceding segment

As we have seen, the antecedents of more implicit co-referring expressions, like pronouns, are very of-

ten in the same segment as the expressions. The antecedents of less explicit co-referring expressions, like deictic and definite NPs, are mostly in the immediately preceding segment, but also often in the same segment. About 50% are 2 utterances away, 20% between 3 and 5, although we find distances up to 41 utterances for definite NPs.

However, there is a small number (11) of cases in which the antecedents are found across the boundaries of the immediately preceding segment. This poses a challenge to systems since the context needed for recovering these antecedent is not as local. The following example is a case of split antecedents. The antecedent of the elided NP is to be found across the two immediately preceding questions. Moreover, as you can see, the Wizard is not sure about how to interpret the missing argument, which can be because of the split antecedents, but also because of the amount of time passed, and/or because one of the answers is still missing, that is, more than one segment is open at the same time.

(11) US: Which are the webpages for *European Joint Conferences on Theory and Practice of Software and International Conference on Linguistic Evidence*?

LT: Please wait... (waiting time)

US: Which are the webpages for *International Joint Conference on Neural Networks and Translating and the Computer 27*?

LT: <http://www.complang.ac>, ... (1st answer)

US: Up to which date is it possible to send a paper, an abstract [J]?

LT: <http://uwb.edu/ijcnn05/>, ... (2nd answer)

LT: For which conference?

US: For *all of the conferences I got the webpages*.

In the following example the antecedent of the definite NP is also to be found beyond the boundaries of the immediately preceding segment.

(12) US: What is the homepage of *the project*?

LT: <http://dip.semanticweb.org>

USER: What is the email address of Christoph Bussler?

LT: The database does not contain this information.

US: Where does *the project* take place?

Here the user asks about the email address of a person who was previously introduced in the discourse as the coordinator of the project under discussion and then keeps on referring to the project with a definite NP. The intervening question is somehow related to the project, but not directly. There is a topic shift, as defined by Chai and Jin (2004), where the main topic becomes an entity related to the entity the preceding question was about. However, this topic shift is only at a very local level, since the dialogue participants keep on speaking about the project, that is, the topic at a more general level keeps on being the same. We can speak here of thematic nesting, since the second question is about an entity introduced in relation to the entity in focus of attention in the first question, and the third question is again about the same entity as the first. The project has not completely left the focus, but has remained in secondary focus during the second segment, to become again the main focus in the third segment. It seems that as long as the entity to which the focus of attention has shifted is related to the entity previously in focus of attention, the latter still also remains within the focus of attention.

5 Conclusions

The possibility of using contextual phenomena is given by certain types of thematic relatedness - especially *theme-entity* and *theme*, for co-reference and bridging, and *refinement*, *theme-entity* and *theme-property*, for fragments -, and contiguity of questions. As we have seen, the immediately preceding segment is in most cases the upper limit of the search space for the last reference to the entity, or the elided material in fragments. The directions of the search for antecedents, however, can vary depending on the phenomena, since for more implicit referring expressions antecedents are usually to be found in the same segment, while for less implicit referring expressions they are to be found in the preceding one.

These data are in accordance with what Ahrenberg et al. (1990) predict in their model. Just to consider the immediately preceding segment as the upper limit of the search space for antecedents is enough and, thus, no tracking of thematic relations is needed to resolve discourse phenomena. How-

ever, there are occurrences of more explicit types of co-reference expressions, where the antecedent is beyond the immediately preceding segment. As we have observed, in these cases the intervening segment/s shift the focus of attention to an entity (maybe provided in some previous answer) closely related to the one in focus of attention in the preceding segment. It seems that as long as this relation exists, even if there are many segments in between⁸, the first entity remains in focus of attention and can be referred to by an implicit deictic or definite NP without any additional retrieval cue. We can speak of thematic nesting of segments, which seems to be analogous to the intentional structure in task-oriented dialogues as in (Grosz and Sidner, 1986), also allowing for reference with implicit devices to entities in the superordinate segments after the subordinated ones have been closed. It seems, thus, that thematic structure, like the discourse goals, also imposes structure on the discourse.

These cases, although not numerous, suggest that a more complex discourse structure is needed for QA interactions than one simply based on the discourse goals. The local context is given by the discourse segments, which are determined by the discourse goals, but a less local context may encompass several segments. As we have seen, reference with implicit devices to entities in the less local context is still possible. What seems to determine this less local context is a unique theme, about which all the segments encompassed by the context directly or indirectly are. So, although it does not seem necessary to track all the thematic transitions between the segments, it seems necessary to categorize the segments as being about a particular more global theme.

In a system like the one we simulated, having specific tasks in mind and querying structured data, a possible approach to model this extended context, or focus of attention, would be in terms of frames. Every time a new entity is addressed a new frame is activated. The frame encompasses the entity itself and the properties holding of it and other entities, as well as those entities. This would already allow us to successfully resolve bridging and fragments with a partial source. If the focus of atten-

⁸We found up to five intervening segments, one of them being a subsegment.

tion then shifts to one of the related entities, the user demanding particular information about it, then its frame is activated, but the previous frame also remains somehow active, although to a lesser degree. As long as there is a connection between the entities being talked about and a frame is not explicitly closed, by switching to speak about a different entity of the same class, for example, frames remain somehow active and implicit references will be accommodated within the activation scope.

In principle, the closer the relation to the entity currently in focus, the higher the degree of activation of the related entities. Yet, there may be cases of ambiguity, where only inferences about the goals of the user may help to resolve the reference, as in (13):

(13) US: How is the contact for that project?

LT: daelem@uia.ua.ac.be

US: What is the institute?

LT: Centrum voor Nederlandse Taal en Spraak.

US: Homepage?

Here the property "Homepage" could be asked about the institution or the project, the institution being more active. However, the Wizard interpreted it as referring to the project without hesitation because she knew that subjects were interested in projects, not in organizations. In order to resolve the ambiguity, we would need a system customized for tasks or make inferences about the goals of the users based on the kind of information they've been asking for. Determining at which level of nesting some expression has to be interpreted may involve plan recognition.

However, for open domain systems not having a knowledge-base with structured data it may be much more difficult to keep track of the focus of attention beyond the strictly local context. For other kinds of interactions which don't have such a structured nature as our tasks, this may also be the case. For example, in the information browsing tasks in (Kato et al., 2004), there is not a global topic encompassing the whole interaction, but the information needs of the user are given by the information he is encountering as the interaction proceeds, that is, he is browsing the information in a free way, without having particular goals or particular pieces of information he wants to obtain in mind. In such cases it may be difficult to determine how long frames are

active if the nesting goes very far, as well as making any inferences about the user's plans. However, it might also be the case, that in that kind of interactions no implicit referring expressions are used beyond the segmental level, because there is no such an extended context. In order to find out, a study with interactive data should be carried out.

Acknowledgements

The research reported here has been conducted in the projects QUETAL and COLLATE II funded by the German Ministry for Education and Research, grants no. 01IWC02 and 01INC02, respectively. We are also grateful to Bonnie Webber for her helpful comments on the contents of this paper.

References

- Ahrenberg Lars, Dahlbäck Nils and Arne Jönsson 1990. *Discourse representation and discourse management for natural language interfaces*. Proceeding of the Second Nordic Conference on Text Comprehension in Man and Machine, Täby, Sweden, 1990.
- Jaime G. Carbonell. 1983. *Discourse pragmatics and ellipsis resolution in task-oriented natural language interfaces*. Proceedings of the 21st annual meeting on Association for Computational Linguistics, Cambridge, Massachusetts, 1983
- Chai Joyce Y. and Ron Jin. 2004. *Discourse Status for Context Questions*. HLT-NAACL 2004 Workshop on Pragmatics in Question Answering (HLT-NAACL 2004) Boston, MA, USA, May 3-7, 2004
- Dahlbäck Nils and Arne Jönsson. 1989. *Empirical Studies of Discourse Representations for Natural Language Interfaces*. Proceedings of the Fourth Conference of the European Chapter of the ACL (EACL'89), Manchester.
- Grosz Barbara and Candace Sidner. 1986. *Attention, Intention and the Structure of Discourse*. Computational Linguistics 12(3): 175-204.
- Kato Tsuneaki, Fukumoto Junichi, Masui Fumito and Noriko Kando. 2004. *Handling Information Access Dialogue through QA Technologies - A novel challenge for open-domain question answering*. HLT-NAACL 2004 Workshop on Pragmatics in Question Answering (HLT-NAACL 2004) Boston, MA, USA, May 3-7, 2004
- Sun Mingyu and Joycie J. Chai. 2006. *Towards Intelligent QA Interfaces: Discourse Processing for Context Questions*. International Conference on Intelligent User Interfaces, Sydney, Australia, January 2006