Lexicalisation Strategies in Cooperative Question-Answering Systems

Farah BENAMARA, Patrick SAINT-DIZIER

IRIT, 118 route de Narbonne
31062 TOULOUSE Cedex, FRANCE
benamara, stdizier@irit.fr

Abstract

In this project note, we present the main features of lexicalisation strategies deployed by humans in question-answering (QA) tasks. We then show how these can be reproduced in automated QA systems, in particular in Intelligent Cooperative Question-Answering Systems.

1 Motivations

The framework of Advanced Question Answering (QA) systems, including cooperative QA answering, as described in a recent road map (Voorees 2003), raises new challenges for language generation since answers are often produced via dedicated inference mechanisms operating over various knowledge sources, including conceptual ontologies. Producing well-formed and informative responses gives a whole new insight to language generation, and in particular to the difficult problem of lexicalisation.

For example, the question:

A bungalow in Megève for 15 persons ?

has the response:

There are no bungalows in Megève, we can propose you 2 close-by chalets or a country cottage for 15 persons in Megève, or a 15 person bungalow in les Houches (8kms) or in Acacias (12kms),

where closely related alternatives are proposed on: (1) sister concepts of bungalow (the question focus): a country cottage or 2 close-by chalets (assuming a chalet capacity does not exceed 10 persons) in the same village or (2) sister concepts of the Megève village (the strongest constraint on the focus): close-by villages where bungalows meet the 15 person constraint. The response involves evaluating village proximity and sorting responses, e.g. by increasing distance from Megève. The first part of the response *There are no bungalows in Megève* is the **direct response**, which corrects the user false presupposition, while the remainder of the response reflects the **cooperative know-how** of the responder. Cooperative know-how involves several forms of responses that include relaxations, intensional

calculus, expression of restrictions, of warnings and conditional responses (Benamara et al. 2004).

This simple example shows that, if a direct response cannot be found, several forms of knowledge and reasoning schemas need to be used and that the NL form of the response requires adequate and subtle lexicalisations, often directed by the reasoning schemas.

Lexicalisation is the operation that associates a word or an expression to a concept. It is a major parameter in response production (Reiter et Dale, 1997), (a good synthesis can be found in (Cahill, 1999)). Lexicalisation is often decomposed into two different stages: **lexical choice**, which occurs during content determination, where the lexical term chosen, which may still be underspecified, is dependent on reasoning procedures, the knowledge base contents, and grammatical constraints; and **lexical variation** which is the choice of a particular word or form among possible synonyms or paraphrases. Lexical variation occurs in the surface realizer, and may have some pragmatic connotations, for example an implicit evaluation via the language level chosen (e.g. argotic entails low evaluation).

In this project note, we show the different facets of lexicalisation in advanced QA, where the two stages, lexical choice and variation, play crucial and very distinct roles in the different parts of the response. Our model is based on a detailed analysis of QA pairs, and on annotation tasks, from which a model emerged, suitable for NLG.

2 The WEBCOOP project and its environment

This research is part of the WEBCOOP system (Benamara et al. 2004), that makes a heavy use of a rich domain ontology, which is a concept ontology structured by means of several semantic relations (Cruse 1986). To better understand the perspective in which this work has been carried out, let us say a few words about the WEBCOOP project, a system designed to provide cooperative answers to queries on the Web.

2.1 Typology of the domain ontology

The ontology on which our project is based is basically a conceptual ontology where nodes are associated with concept lexicalizations and essential properties, as it is now often the case in restricted domain ontologies (<http://www.daml.org/ontologies/>). In our implementation, each node is represented by the predicate :

```
onto-node(concept, lex, prop)
where concept has properties prop, represented
as attribute-value pairs, and lexicalisations lex.
Most lexicalisations are entries in the lexicon,
where morphological and grammatical aspects are
described. For example, for hotel, we have (coded
in Prolog):
```

```
onto-node(hotel,
[[hotel], [residence, hoteliere]],
[night-rate= X: integer ,
nb-of-rooms= N: integer, facilities:
List]).
```

There are several well-designed public domain ontologies on the net. Our ontology is a synthesis of two existing French ontologies, that we customized: TourinFrance (www.tourinfrance.net) and the bilingual (French and English) thesaurus of tourism and leisure activities (www.iztg.hr/indokibiblioteka/THESAUR.PDF) which includes 2800 French terms. We manually integrated these ontologies in WEBCOOP (Benamara et al. 2004a) and added ourselves most properties by hand. We also revised slightly the ontology by removing concepts which are either too specific (i.e. too low level), such as some basic aspects of ecology or rarely considered, as e.g. the economy of tourism. We also removed quite surprising classifications such as *sanatorium* under *tourist accommodation*. We finally reorganized some concept hierarchies, so that they 'look' more intuitive for a large public. Finally, we found that some hierarchies are a little bit odd, for example, we found at the same level *accommodation capacity* and *tourist accommodation* whereas, in our case, we consider that *capacity* is a property of the concept *tourist accommodation*.

We have, at the moment, 1000 concepts in our tourism ontology which describe the different aspects of accommodation and transportation and a few other satellite elements (geography, health, immigration). Besides the traditional 'isa' relation, we also coded the 'part-of' relation and the opposition. We also have a few proportional series in order to encode graduality. In the tourism domain there is a reasonable number of part-of pairs, and a small number of opposites and proportional series. Syn-

onymy is encoded via the list of lexicalizations. Encoding these relations is straightforwardly realized by introducing facts, which operate at the concept level, not at the lexicalization one:

```
part-of(journey, trip).
```

As shall be seen below, this ontology is further annotated for lexicalization tasks.

Finally, in (Benamara et al. 2004), we define a conceptual metrics that determines the distance between two concepts, evaluated in terms of distance in the hierarchy between the two concepts and property differences (nature of properties and values associated).

2.2 Outline of WEBCOOP

Following Grice's maxims (Grice, 1975) and related works, e.g. (Searle, 1975), in the early 1990s, a number of forms of cooperative responses were identified. Most of the efforts in these studies and systems focussed on the foundations and on the implementation of reasoning procedures (Gal, 1988), (Minock et al., 1996), while little attention was paid to question analysis and NL response generation. An overview of these systems can be found in (Gasterland et al., 1994) and in (Webber et al., 2002), based on works by (Hendrix et al., 1978), (Kaplan, 1982), (Mays et al., 1982), among others. These systems include e.g. the identification of false presuppositions and various types of misunderstandings found in questions. They also include reasoning schemas based e.g. on constant relaxation to provide approximate or alternative, but relevant, answers when the direct question has no response. Intensional reasoning schemas can also be used to generalize over lists of basic responses or to construct summaries.

The framework of Advanced Reasoning for Question Answering (QA) systems, as described in a recent road map, raises new challenges since answers can no longer be only directly extracted from texts (as in TREC) or databases, but require the use of a domain knowledge base, including a conceptual ontology, and dedicated inference mechanisms. Such a perspective, obviously, reinforces and gives a whole new insight to cooperative answering.

In WEBCOOP, user questions may range from keywords to comprehensive natural language expressions. Their parse produces a semantic logical representation that includes the question conceptual category, the question focus and the question contents. Responses are structured in two parts. The first part contains explanation elements in natural language. It is a first level of cooperativity that reports user misconceptions, if any, in rela-

tion with the domain knowledge. The second part is the most important and the most original. It reflects the know-how of the cooperative system, going beyond the cooperative statements given in the first part. It is based on intensional description techniques and on intelligent relaxation procedures going beyond classical generalization methods used in AI. This component also includes additional dedicated cooperative rules that make a thorough use of the domain ontology and of general knowledge. Responses provided to users are built in web style, by integrating natural language generation (NLG) techniques with hypertext links to produce "dynamic" responses. Hyperlinks are dynamically created at generation time. This leaves up to the user the high-level planning tasks inherent to NLG and improves readability and information access. The standard generation difficulties (lexicalisation, aggregation, argumentation) remain crucial to generate cooperative responses, but their web-style greatly reduces the overall complexity.

The goal of the present contribution is to clarify and model the lexicalization aspects, which are, in fact, in close dependence with the aggregation functions.

3 Investigating lexicalisations from cooperative QA corpora

To carry out our study, we considered three typical sources of cooperative discourses: Frequently Asked Questions (FAQ), Forums, and email question-answer pairs (EQAP), these latter obtained by sending ourselves emails to relevant services (e.g. for tourism: tourist offices, airlines, hotels). Our study was carried out on 670 cooperative question-answer pairs. We have about 50% pairs coming from FAQ, 25% from Forums and 25% from EQAP. The domains considered are basically large-public applications: tourism (45%), health (22%), sport, shopping and education (19%). In all these corpora, no user model is assumed, and there is no dialogue: QA pairs are isolated, with no context, in a way similar to web search engine querying.

4 Lexicalisation strategies deployed by humans

From the analysis of these QA pairs, we have identified the main lexicalisation strategies presented below. Between parentheses are introduced XML-type labels, used to annotate our corpora.

- Identity (**ID**): the term in the question is kept unchanged in the response,

- Quasi-synonyms (**SYN**): under this term, we observe a large spectrum of variations:

SYN(a) terms which are almost equivalent (*have/possess*),

SYN(b) usages with a specific orientation or connotation, such as the use of commercially oriented terms to attract customers or of more technical terms, possibly with slight nuances, on the part of the responder (*refund / financial compensation, breathing stops / respiration pauses*), quite typical of financial or health corpora where the responder is often a professional,

SYN(c) transcategorial variations or short paraphrases (*international → from the entire world, from abroad, from all the countries in the world*). This latter subclass is much more difficult to characterize, it often indicates a form of insistence or stress, or the desire to make information more explicit (e.g. *international* is really *from all the countries in the world*),

- Opposites (**OPP**): used on a limited scale to account e.g. for the differences in point of view between the questioner and the responder (e.g. *lend / borrow, pay/charge, stay/leave*) (Cruse 1986).

- Subtypes (**SUB**): often subtypes in the ontology (*accommodation → hotel, motel, bungalow, country cottage, etc.*) or enumerations of basic terms (e.g. different types of credit cards). We have the following main subtypes:

SUB(a) lower-level elements in an ontology (possibly terminal elements),

SUB(b) focuses on sub-events or sub-activities of the larger event lexicalised in the question, or (simple) procedure description instead of the procedure name, and

SUB(c) subtyping via modification (*official guide, mountain guide*).

- Sisters (**SIS**) are used when the response introduces minor superficial corrections to the terms used in the question (not as important as false presuppositions), often related to 'lexical approximations' (e.g. *bungalow* used instead of *hut, cabin* in a camping), these should not be confused with SYN(b). SIS is also useful in negative responses, where the correct response is a sister of the concept used in the query (*is my metro ticket valid on the bus ? no, you need to buy a bus ticket*).

- Generalizations (**GEN**): their main use is in intensional answers, where generalizations are

realized to make the answer more compact. Generalizations can also be used e.g. to insist on a certain facet of the term (*castle / building*), or to provide a certain form of help (*where is the Dominican Republic ? → in the Antilles*). It is possible, in addition, to indicate the generalization level w.r.t. the ontology.

- Metonymies (**MET**): are relatively frequent, and focus on a particular ‘facet’ of the term used in the question to better emphasize a particular aspect, such as a means or an instrument of an action, or use of make for object, container for containee.
- Fuzzy term interpretation (**FTI**): *close* interpreted as *50 meters*.

These lexical variations are relatively well distributed over the above different categories, roughly as follows (all situations included) in % :

ID	SYN	OPP	SUB	SIS	GEN	MET	FTI
32%	21%	2%	21%	4%	5%	14%	2%

As the reader may note it, almost none of these lexical variations are neutral. They convey a specific connotation, point of view, focus or they may activate a concept property or facet usually in the background.

To have a more systematic analysis of lexicalisation, we manually annotated lexicalisation forms in QA pairs. The examples in Fig. 1, on the next page, illustrate several forms frequently encountered. Terms considered are labelled as T1, T2, etc. (including terms within the focus), a special annotation is reserved for the whole structure identified as the focus (FOC) and op= specifies the variation operation at stake (e.g. SYN(a), SUB(b), etc.). Q1-R1 have several layers of lexicalisations: *guided tour* is a subtype of *visit*, *guided tours at fixed hours 11AM and 3 PM* is a subtype of *visit hours*, while *building* is a generalization of *castle*, necessary to allow for the reference to *town hall*, since *building* is the least upper bound in the ontology for *castle* and *town hall*. Q2-R2 stresses on and makes more explicit *drinkable* in *water from the tap*, while Q3-R3 requires the GEN strategy, which is the simplest way to locate a country or any object, via its relations with others.

5 Enriching the ontology for lexicalisations tasks in NLG

Let us consider again the ontology presented in section 2.1. It needs some refinements so that the lexicalization situations identified can be automated. Let us consider the main cases.

SYN(a) is based on the set of lexicalizations associated with the concept considered. SYN(b) requires lexicalisations to be marked according to their language level, in our case: slang, popular term, standard term (the one chosen by default) and technical term. SYN(c) requires another type of notation: the category of the lexicalisation: word versus expression + category. For example, *possess* is marked technical (contrasted with *have*), and *from the entire world* is identified as an expression of type PP. In terms of coding, the last example is coded as:

```
lex([from,the,entire,world],
standard, expr:pp).
```

Opposites are encoded via a specific set of facts in the ontology at the concept level, same for SUB(b) where subactivities of an activity are considered as parts. Treating SIS is more delicate since all the sisters of a node may not be appropriate. At this level, we use the conceptual metrics advocated in section 2.1, with a quite high threshold, that we parameterize to evaluate its best level. The metrics is highly dependent on the number of properties encoded, but it seems hard to have another strategy, since sampling methods would require too much corpus work. As an example, *hut* and *cabin* are considered as close sisters of *bungalow*, but *chalet* and *country cottage* are more remote. This aspect is crucial when relaxing concepts to find alternative solutions.

Finally, metonymies will not be treated at the moment and FTI are dealt with in a more or less ad’hoc way, using simple and naive geometrical considerations.

6 A lexicalisation model for QA

As shown in Fig. 2, lexicalisation strategies in QA pairs occur at five levels: **L1**: between the query and the direct response, **L2**: between the question and the cooperative know-how part, **L3**: within the direct response (insistence, paraphrases, exclusions, etc.), **L4**: between the direct response and the know-how part, and, finally, **L5**: within the know-how part, which may be composed of several components, as shown in (Benamara et al. 2004b). It seems that L1, L2, L3, L4 and L5 are organized hierarchically in the lexicalisation process.

6.1 The relation L1

Let us concentrate on the question focus, which is the term the most sensitive to lexical variation (variations on other terms are essentially synonyms). Let us first consider **L1**, the lexicalisation strategies between the question and the direct response. The

Q1: What are the < FOC > < T2 > < T3 > visit < /T3 > hours < /T2 > of the < T1 > castle ? < /T1 > < /FOC >
R1: This < T1 op = GEN > building < /T1 > being the town hall, there are only < T2 op = SUB > < T3 op = SUB > guided tours < /T3 > at fixed hours: 11 AM and 3 PM < /T2 >.
Q2: Is water < FOC > drinkable < /FOC > in the Dominican Republic ?
R2: No, you must not < FOC op = SYN(c) > drink water from the tap < /FOC >, always water from bottles.
Q3: < FOC > Where < /FOC > is the Dominican Republic ?
R3: It is in the very heart of < FOC op = GEN > the Caribbean < /FOC >, it is the second largest < FOC op = GEN > island of the Antilles < /FOC >.

Figure 1: Lexicalisation annotations

Figure 2: Lexicalisation relations in a QA pair

main criterion turns out to be the ‘conceptual’ overlap, observed in the domain ontology, between the response elements and the question focus. This is determined by comparing the subtree ST1 in the domain ontology whose root is the concept associated with the question focus, with the subtree ST2 associated with the response, or, in case of an enumeration, the smallest subtree that includes all the solutions. We basically have four situations:

- Focus adequate w.r.t. the response: both subtrees coincide, or ST2 is a proper subtree of ST1, (Q: *what kind of tourist accommodation in Chamonix?* R: *we have all kinds of hotels, bungalows, chalets and campings*),
- Focus too large: ST2 is only a portion of a proper subtree of ST1 (Q: *what kind of tourist accommodation in La Trairie?* *we have only cabins and bungalows*),
- Focus too narrow: ST2 is larger than ST1, i.e. ST2 dominates ST1 or has elements not included into ST1, for example: (Q: *An express train to Marseille at 11.PM ?* R: *We have an express train at 11.15, but also a TGV, which is much faster, at 11.25.*, where there is a generalization over the types of trains,
- Shifts: ST2 is not included into ST1 (but, in some cases, a few solutions can be common) (*what hotel categories ? we only have campsites and bungalows*).

The above situations are categorized in terms of the lexicalisation functions below, percentage per

focus situation is given for L1 only, more pragmatic aspects are briefly commented:

Focus	Lexicalisation Strategies
(1) Adequate, Response positive or negative (40%)	ID, SYN(a) SYN(b) specific orientation SYN(c) for stress SUB(a)(b) for a more refined response via enumeration
(2) Too large (21%)	SUB(c), MET, SUB(a,b) for enumerations
(3) Too restricted (6%)	GEN reference to generic concepts
(4) Fuzzy or needs an interpretation (2%)	FTI, SYN not interpreted
(5) Adequate with conditions (18%)	SUB(a)(b) for enumerations
(6) Minor shifts (13%)	SIS for adjustments

In an NLG system, more schematic than real-life corpora, the above strategies allow for the production of adequate lexicalisations for the following situations, which seem to cover all the cases we found for yes/no and entity (what is, who) questions: illustration: Q: *what kinds of hotels in Megève?*

- direct response, simple or complex, (ID, SYN), *we have all kinds of hotels in Megève*.
- response with restrictions (SYN, ID + SUB, GEN + SUB), *We have all kinds of tourist accommodation except 4* hotels and huts*, or simple adaptations (SIS),
- response scoping on precise subtopics or subtypes (SUB), response can be negative on the query focus and then positive on a subtype: *We do not have hotels, but comfortable, high standard, chalets*,

- response with some form of generalization (as in example Q1), *We have all types of tourist accomodation, from the cheapest to the fanciest* (GEN, least upper bound),
- response with interpretation of fuzzy term (FTI),
- response with conditions, or case structures (SUB), *we have 2* hotels at 30-60 Euros, 3* hotels at 50-80 Euros, etc.*
- response with forms of stress, insistence, explication via enumerations or paraphrases (SYN, SUB, OPP, GEN), *we have all types of hotels: 2*, 3*, 4* and even 5* hotels.*

The implementation of L1 is functionally quite straightforward. The reasoning component produces a logical formula, from which the NLG functions (lexicalisation, but also aggregation and microplanning) operate. As indicated in the above array, the generator evaluates the relative overlap between the question focus and the response. It can then activate, for the terms to generate, the relevant lexicalisation function(s). The GEN function is quite specific since it has also some stylistic motivations, it is described in depth in (Benamara 2004). There are, at this level, several choices, that we manage via preferences given a priori since we do not have any user model. Stress and various additional enumerations have essentially a pragmatic motivation, as indicated above. At the implementation level, these can just be parameterized.

6.2 The relations L2, L3 and L4

Within the direct response, the most direct lexicalisation forms (L3, Fig. 2), can be paired with categories like OPP, SUB, FTI or MET. OPP is used within the response to reinforce some terms by saying what they are not. MET is in general a form of stress or focus on a particular point. MET is produced via the concept properties of the ontology. SUB(a)(b) are also used as forms of elaborations, to make more explicit a term (e.g. *we accept all forms of payments: cheques, credit cards, cash, money orders, etc.*), or to give a precision, possibly using FTI. Such lexicalisations have a strong pragmatic basis. However, in NLG, the use of OPP or of SYN(a)(b) can be motivated either by the user model, or, if it does not exist, by the need of some elaboration if some concepts in the response are complex or quite remote from the original question focus. The implementation relies on the conceptual metrics, also

used for concept relaxation (Benamara et al 2004a) and also on manually marked really complex concepts in the ontology. The lexicalization programme is simply able to propose additional lexicalizations, structured by the aggregation procedure, e.g. as a list of SUB or OPP.

In the know-how part of the cooperative response, lexicalisation strategies (L2) are more rigid, since they are guided by the reasoning procedures. We observe a large number of relaxations, involving lexicalisations based on SUB, SIS or GEN, depending on how the ontology is traversed by relaxation. Several intensional operations are also developed to make the response more compact, involving lexicalisations based on GEN (Benamara 2004). Lexicalization choices totally depend on the reasoning procedures and are just direct language realizations of the concepts introduced by relaxation or generalization inference schemas.

There are lexicalisation relations between the direct response and the know-how part (L4, which has some relations with L2) in particular when the direct response is negative. The know how part involves realisations which are in:

- GEN, SUB, SIS or FTI relation with the term used in the direct response when relaxation or intentional calculus is used in the know how part,
- SYN or SUB relation when the know-how part is a warning or a restriction. Pronominal or other forms of references are possible at this level (*yes, we do have cabins, but they can only be rent on a weekly basis*).

6.3 The relation L5

The relation L5 has a specific status, it covers a large number of situations, we therefore simply evoke prototypical situations here. Its motivation comes from the fact that the know how part of a response is composed, in general, of several segments, called meaningful units. As explained in detail in (Benamara et al. 2004b), the know-how part is in general composed of a response elaboration (related to constraint relaxation, or any other type of inference) followed by a number of additional information under the form of justifications, warnings, restrictions, counter-proposals, etc.

For example, in the following example, we have several segments:

Q: *How can I get to the Borme castle ?*

Direct response + elaboration: *You must take the GR90 from the old castle:*

additional information - precision: *walking distance: 30 minutes.*

warning: *There is no possibility to get there by car.*

In this example, the warning segment uses a kind of opposite of *walking*, namely *driving*, to warn the user e.g. about the lack of alternative.

The following example has a know-how part based on relaxation, and then proposes an alternative to cinemas by means of a sister concept:

Q: *Is there a cinema in Borme ?*

Direct response: *No*,

Know-how (relaxation on the localization): *the closest cinema is at Londes (8 kms) or at Hyeres, Cinema Olbia (at 20 kms).*

Additional information (counter-proposal on the activity) : *however, there is a nice open air theater in Borme.*

In this latter example, we see that, again, the additional information segment introduces a sister concept of cinema.

Implementing L5 is largely dependent on the reasoning procedures, on the knowledge found in the database, and on the way the general form of the response is planned. The role of lexicalisation here is essentially choosing the appropriate terms, with the appropriate level in the ontology.

7 Perspective

This study shows several original lexicalisation strategies in advanced, cooperative QA forms. Lexical choice, realized in the content determination phase in conjunction with the cooperative inference schemas, mainly involves the SIS, GEN, SUB, FTI and ID functions, whereas lexical variation, realized at the level of the surface realizer, along the five relations L1, L2, L3, L4 and L5, mainly involves SYN, OPP and MET, characterized by the production of lists of potential candidate words. This latter level interacts also with grammatical considerations such as aggregation (Reiter and Dale 1997) or the use of alternations.

These lexicalisations require a quite advanced domain model, necessary for the development of cooperative functions, in particular a rich ontology. These functions have been tested in the WEBCOOP prototype, that cooperatively responds to queries on the WEB, on the tourism domain (accommodation and transportation).

This study can be applied to open generation, on a purely conceptual representation basis, or to more restricted forms of generation, such as underspecified templates with calls to lexicalisation.

Acknowledgements We thank the CNRS GDR TCAN and Véronique Moriceau for providing us corpora and for helpful discussions. We also thank the three reviewers for their useful comments.

References

- Benamara, F, Saint Dizier, P., *Advanced Relaxation for Cooperative Question Answering*, in New Directions in Question Answering, Chapter 21. Mark T. Maybury (ed.), AAAI/MIT Press, 2004a.
- Benamara, F., Moriceau, V., Saint-Dizier, P., *COOPML: Towards Annotating Cooperative Discourse*, in Workshop on discourse annotation, ACL'04, Barcelona, July 2004b.
- Benamara, F., AAAI.
- Cahill, L., Lexicalisation in applied NLG systems, Research report, ITRI-99-04, 1999.
- Cruse, A., *Lexical Semantics*, CUP, 1986.
- Gal, A., *Cooperative Responses in Deductive Databases*, PhD Thesis, Univ. of Maryland, 1988.
- Gaasterland, T., Godfrey, P., Minker, J., *An Overview of Cooperative Answering*, Papers in non-standard queries and non-standard answers, Clarendon Press, Oxford, 1994.
- Grice, H., *Logic and Conversation*, in Cole and Morgan (eds), Syntax and Semantics, Academic Press, 1975.
- Hendrix, G., Sacerdoti, E., Sagalowicz, D., Slocum, J., Developing a Natural Language Interface to Complex Data, *ACM transactions on database systems*, 3(2), 1978.
- Kaplan, J., Cooperative Responses from a Portable Natural Language Query System, in M. Brady and R. Berwick (ed), *Computational Models of Discourse*, 167-208, MIT Press, 1982.
- Lehnert, W., *The Process of Question Answering: a Computer Simulation of Cognition*, Lawrence Erlbaum, 1978.
- Mays, E., Joshi, A., Webber, B., Taking the Initiative in Natural Language Database Interactions: Monitoring as Response, *EACL'82*, Orsay, France, 1982.
- Minock M, Chu W, Yang H, Chiang K, Chow, G and Larson, C, *CoBase: A Scalable and Extensible Cooperative Information System*. Journal of Intelligent Information Systems, volume 6, number 2/3, pp : 223-259, 1996.
- Reiter, R., Dale, R., *Building Applied Natural Language Generation Systems*, Journal of Natural Language Engineering, volume 3, number 1, pp:57-87, 1997.
- Searle, J., *Indirect Speech Acts*, in Cole and Morgan (eds), Syntax and Semantics III, Academic Press, 1975.
- Voorees, E. M., *Overview of the TREC 2002 Question Answering Track*. Proceedings of TREC-11, NIST, 2003.