THE TEXT SYSTEM FOR NATURAL LANGUAGE GENERATION:
AN OVERVIEW*

Kathleen R. McKeown
Dept. of Computer & Information Science
The Moore School
University of Pennsylvania
Philadelphia, Pa. 19104

ABSTRACT

Computer-based generation of natural language requires consideration of two different types of problems: 1) determining the content and textual shape of what is to be said, and 2) transforming that message into English. A computational solution to the problems of deciding what to say and how to organize it effectively is proposed that relies on an interaction between structural and semantic processes. Schemas, which encode aspects of discourse structure, are used to guide the generation process. A focusing mechanism monitors the use of the schemas, providing constraints on what can be said at any point. These mechanisms have been implemented as part of a generation method within the context of a natural language database system, addressing the specific problem of responding to questions about database structure.

1.0 INTRODUCTION

Deciding what to say and how to organize it effectively are two issues of particular importance to the generation of natural language text. In the past, researchers have concentrated on local issues concerning the syntactic and lexical choices involved in transforming a pre-determined message into natural language. The research described here emphasizes a computational solution to the more global problems of determining the content and textual shape of what is to be said. More specifically, my goals have been the development and application of principles of discourse structure, discourse coherency, and relevancy criterion to the computer generation of text. These principles have been realized in the TEXT system, reported on in this paper.

The main features of the generation method used in TEXT include 1) an ability to select relevant information, 2) a system for pairing rhetorical techniques (such as analogy) with discourse purposes (such as defining terms) and 3) a focusing mechanism. Rhetorical techniques, which encode aspects of discourse structure, guide the selection of information for inclusion in the text from a relevant knowledge pool - a subset of

the knowledge base which contains information relevant to the discourse purpose. The focusing mechanism helps maintain discourse coherency. It aids in the organization of the message by constraining the selection of information to be talked about next to that which ties in with the previous discourse in an appropriate way. These processes are described in more detail after setting out the framework of the system.

2.0 APPLICATION

In order to test generation principles, the TEXT system was developed as part of a natural language interface to a database system, addressing the specific problem of generating answers to questions about database structure. Three classes of questions have been considered: questions about information available in the database, requests for definitions, and questions about the differences between database entities [MCKEOWN 80]. In this context, input questions provide the initial motivation for speaking.

Although the specific application of answering questions about database structure was used primarily for testing principles about text generation, it is a feature that many users of such systems would like. Several experiments ([MALHOTRA 75], [TENNANT 79]) have shown that users often ask questions to familiarize themselves with the database structure before proceeding to make requests about the database contents. The three classes of questions considered for this system were among those shown to be needed in a natural language database system.

Implementation of the TEXT system for natural language generation used a portion of the Office of Naval Research (ONR) database containing information about vehicles and destructive devices. Some examples of questions that can be asked of the system include:

- > What is a frigate?
- > What do you know about submarines?
- > What is the difference between a whisky and a kitty hawk?

*This work was partially supported by National Science Foundation grant #MCS81-07290.

The kind of generation of which the system is capable is illustrated by the response it generates to question (A) below.

A) What kind of data do you have?

All entities in the ONR database have DB attributes REMARKS. There are 2 types of entities in the ONR database: destructive devices and vehicles. The vehicle has DB attributes that provide information on SPEED-INDICES and TRAVEL-MEANS. The destructive device has DB attributes that provide information on LETHAL-INDICES.

TEXT does not itself contain a facility for interpreting a user's questions. Questions must be phrased using a simple functional notation (shown below) which corresponds to the types of questions that can be asked. It is assumed that a component could be built to perform this type of task and that the decisions it must make would not affect the performance of the generation system.

1. (definition <e>)
2. (information <e>)
3. (difference <e1> <e2>)

where <e>, <e1>, <e2> represent entities in the database.

3.0 SYSTEM OVERVIEW

In answering a question about database structure, TEXT identifies those rhetorical techniques that could be used for presenting an appropriate answer. On the basis of the input question, semantic processes produce a relevant knowledge pool. A characterization of the information in this pool is then used to select a single partially ordered set of rhetorical techniques from the various possibilities. A formal representation of the answer (called a "message") is constructed by selecting propositions from the relevant knowledge pool which match the rhetorical techniques in the given set. The focusing mechanism monitors the matching process; where there are choices for what to say next (i.e. - either alternative techniques are possible or a single technique matches several propositions in the knowledge pool), the focusing mechanism selects that proposition which ties in most closely with the previous discourse. Once the message has been constructed, the system passes the message to a tactical component [BOSSIE 81] which uses a functional grammar [KAY 79] to translate the message into English.

4.0 KNOWLEDGE BASE

Answering questions about the structure of the database requires access to a high-level description of the classes of objects in the database, their properties, and the relationships between them. The knowledge base used for the TEXT system is a standard database model which draws primarily from representations developed by Chen [CHEN 76], the Smiths [SMITH and SMITH 77], Schubert [SCHUBERT et. al. 79], and Lee and Gerritsen [LEE and GERRITSEN 78]. The main features of TEXT's knowledge base are entities, relations, attributes, a generalization hierarchy, a topic hierarchy, distinguishing descriptive attributes, supporting database attributes, and based database attributes.

Entities, relations, and attributes are based on the Chen entity-relationship model. A generalization hierarchy on entities [SMITH and SMITH 77], [LEE and GERRITSEN 78], and a topic hierarchy on attributes [SCHUBERT et. al. 79] are also used. In the topic hierarchy, attributes such as MAXIMUM SPEED, MINIMUM SPEED, and ECONOMIC SPEED are generalized as SPEED INDICES. In the generalization hierarchy, entities such as SHIP and SUBMARINE are generalized as WATER-GOING VEHICLE. The generalization hierarchy includes both generalizations of entities for which physical records exist in the database (database entity classes) and sub-types of these entities. The sub-types were generated automatically by a system developed by McCoy [MCCOY 82].

An additional feature of the knowledge base represents the basis for each split in the hierarchy [LEE and GERRITSEN 78]. For generalizations of the database entity classes, partitions are made on the basis of different attributes possessed, termed supporting db attributes. For sub-types of the database entity classes, partitions are made on the basis of different values possessed for given, shared attributes, termed based db attributes. Additional descriptive information that distinguishes sub-classes of an entity are captured in distinguishing descriptive attributes (DDAs). For generalizations of the database entity classes, such DDAs capture real-world characteristics of the entities. Figure 1 shows the DDAs and supporting db attributes for two generalizations. (See [MCCOY 82] for discussion of information associated with sub-types of database entity classes).

FIGURE 1 DDAs and supporting db attributes

5.0 SELECTING RELEVANT INFORMATION

The first step in answering a question is to circumscribe a subset of the knowledge base containing that information which is relevant to the question. This then provides limits on what information need be considered when deciding what to say. All information that might be relevant to the answer is included in the partition, but all information in the partition need not be included in the answer. The partitioned subset is called the relevant knowledge pool. It is similar to what Grosz has called "global focus" [GROSZ 77] since its contents are focused throughout the course of an answer.

The relevant knowledge pool is constructed by a fairly simple process. For requests for definitions or available information, the area around the questioned object containing the information immediately associated with the entity (e.g. its superordinates, sub-types, and attributes) is circumscribed and partitioned from the remaining knowledge base. For questions about the difference between entities, the information included in the relevant knowledge pool depends on how close in the generalization hierarchy the two entities are. For entities that are very similar, detailed attributive information is included. For entities that are very different, only generic class information is included. A combination of this information is included for entities falling between these two extremes. (See [MCKEOWN 82] for further details).

6.0 RHETORICAL PREDICATES

Rhetorical predicates are the means which a speaker has for describing information. They characterize the different types of predicating acts s/he may use and delineate the structural

relation between propositions in a text. Some examples are "analogy" (comparison with a familiar object), "constituency" (description of sub-parts or sub-types), and "attributive" (associating properties with an entity or event). Linguistic discussion of such predicates (e.g. [GRIMES 75], [SHEPHERD 26]) indicates that some combinations are preferable to others. Moreover, Grimes claims that predicates are recursive and can be used to identify the organization of text on any level (i.e. - proposition, sentence, paragraph, or longer sequence of text), although he does not show how.

I have examined texts and transcripts and have found that not only are certain combinations of rhetorical techniques more likely than others, certain ones are more appropriate in some discourse situations than others. For example, I found that objects were frequently defined by employing some combination of the following means: (1) identifying an item as a member of some generic class, (2) describing an object's function, attributes, and constituency (either physical or class), (3) making analogies to familiar objects, and (4) providing examples. These techniques were rarely used in random order; for instance, it was common to identify an item as a member of some generic class before providing examples.

In the TEXT system, these types of standard patterns of discourse structure have been captured in schemas associated with explicit discourse purposes. The schemas loosely identify normal patterns of usage. They are not intended to serve as grammars of text. The schema shown below serves the purposes of providing definitions:

Identification Schema

```

identification (class&attribute/function)
[analogy/constituency/attributive]*
[particular-illustration/evidence]+
{amplification/analogy/attributive}
{particular-illustration/evidence}
  
```

Here, "{ }" indicates optionality, "/" indicates alternatives, "+" indicates that the item may appear 1-n times, and "*" indicates that the item may appear 0-n times. The order of the predicates indicates that the normal pattern of definitions is an identifying proposition followed by any number of descriptive predicates. The speaker then provides one or more examples and can optionally close with some additional descriptive information and possibly another example.

TEXT's response to the question "What is a ship?" (shown below) was generated using the identification schema. The sentences are numbered to show the correspondence between each sentence and the predicate it corresponds to in the instantiated schema (the numbers do not occur in the actual output).

(definition SHIP)

Schema selected: identification

- 1) identification
- 2) evidence
- 3) attributive
- 4) particular-illustration

1) A ship is a water-going vehicle that travels on the surface. 2) Its surface-going capabilities are provided by the DB attributes DISPLACEMENT and DRAFT. 3) Other DB attributes of the ship include MAXIMUM SPEED, PROPULSION, FUEL (FUEL CAPACITY and FUEL TYPE), DIMENSIONS, SPEED DEPENDENT RANGE and OFFICIAL NAME. 4) The DOWNES, for example, has MAXIMUM SPEED of 29, PROPULSION of STMTURGRD, FUEL of 810 (FUEL CAPACITY) and BNKR (FUEL TYPE), DIMENSIONS of 25 (DRAFT), 46 (BEAM), and 438 (LENGTH) and SPEED DEPENDENT RANGE of 4200 (ECONOMIC_RANGE) and 2200 (ENDURANCE_RANGE).

Another strategy commonly used in the expository texts examined was to describe an entity or event in terms of its sub-parts or sub-classes. This strategy involves: 1) presenting identificational or attributive information about the entity or event, 2) presenting its sub-parts or sub-classes, 3) discussing attributive or identificational information with optional evidence about each of the sub-classes in turn, and 4) optionally returning to the original entity with additional attributive or analogical information. The constituency schema, shown below, encodes the techniques used in this strategy.

The Constituency Schema

```
attributive/identification (entity)
constituency (entity)
{ attributive/identification
  (sub-class1, sub-class2,...)
  {evidence
    (sub-class1, sub-class2, ...)} }+
{attributive/analogy (entity) }
```

TEXT's response to the question "What do you know about vehicles?" was generated using the constituency schema. It is shown below along with the predicates that were instantiated for the answer.

(information VEHICLE)

Schema selected: constituency

- 1) attributive
- 2) constituency
- 3) attributive
- 4) attributive
- 5) attributive

1) The vehicle has DB attributes that provide information on SPEED INDICES and TRAVEL MEANS. 2) There are 2 types of vehicles in the ONR database: aircraft and

water-going vehicles. 3) The water-going vehicle has DB attributes that provide information on TRAVEL MEANS and WATER GOING OPERATION. 4) The aircraft has DB attributes that provide information on TRAVEL MEANS, FLIGHT RADIUS, CEILING and ROLE. Other DB attributes of the vehicle include FUEL (FUEL CAPACITY and FUEL TYPE) and FLAG.

Two other strategies were identified in the texts examined. These are encoded in the attributive schema, which is used to provide detailed information about a particular aspect of an entity, and the compare and contrast schema, which encodes a strategy for contrasting two entities using a description of their similarities and their differences. For more detail on these strategies, see [MCKEOWN 82].

7.0 USE OF THE SCHEMAS

As noted earlier, an examination of texts revealed that different strategies were used in different situations. In TEXT, this association of technique with discourse purpose is achieved by associating the different schemas with different question-types. For example, if the question involves defining a term, a different set of schemas (and therefore rhetorical techniques) is chosen than if the question involves describing the type of information available in the database.

The identification schema can be used in response to a request for a definition. The purpose of the attributive schema is to provide detailed information about one particular aspect of any concept and it can therefore be used in response to a request for information. In situations where an object or concept can be described in terms of its sub-parts or sub-classes, the constituency schema is used. It may be selected in response to requests for either definitions or information. The compare and contrast schema is used in response to a question about the difference between objects. A summary of the assignment of schemas to question-types is shown in Figure 2.

Schemas used for TEXT

1. identification
-requests for definitions
2. attributive
-requests for available information
3. constituency
-requests for definitions
-requests for available information
4. compare and contrast
-requests about the differences
between objects

FIGURE 2

Once a question has been posed to TEXT, a schema must be selected for the response structure which will then be used to control the decisions involved in deciding what to say when. On the basis of the given question, a set of schemas is selected as possible structures for the response. This set includes those schemas associated with the given question-type (see Figure 2 above). A single schema is selected out of this set on the basis of the information available to answer the question.

For example, in response to requests for definitions, the constituency schema is selected when the relevant knowledge pool contains a rich description of the questioned object's sub-classes and less information about the object itself. When this is not the case, the identification schema is used. The test for what kind of information is available is a relatively simple one. If the questioned object occurs at a higher level in the hierarchy than a pre-determined level, the constituency schema is used. Note that the higher an entity occurs in the hierarchy, the less descriptive information is available about the entity itself. More information is available about its sub-parts since fewer common features are associated with entities higher in the hierarchy.

This type of semantic and structural interaction means that a different schema may be used for answering the same type of question. An earlier example showed that the identification schema was selected by the TEXT system in response to a request for a definition of a ship. In response to a request for a definition of a guided projectile (shown below), the constituency schema is selected since more information is available about the sub-classes of the guided projectile than about the guided projectile itself.

(definition GUIDED)

Schema selected: Constituency

- 1) identification
- 2) constituency
- 3) identification
- 4) identification
- 5) evidence
- 6) evidence
- 7) attributive

1) A guided projectile is a projectile that is self-propelled. 2) There are 2 types of guided projectiles in the ONR database: torpedoes and missiles. 3) The missile has a target location in the air or on the earth's surface. 4) The torpedo has an underwater target location. 5) The missile's target location is indicated by the DB attribute DESCRIPTION and the missile's flight capabilities are provided by the DB attribute ALTITUDE. 6) The torpedo's underwater capabilities are provided by the DB attributes under DEPTH (for example, MAXIMUM OPERATING DEPTH). 7) The guided projectile has DB attributes TIME_TO_TARGET_UNITS, HORZ_RANGE_UNITS and NAME.

Once a schema has been selected, it is filled by matching the predicates it contains against the relevant knowledge pool. The semantics of each predicate define the type of information it can match in the knowledge pool. The semantics defined for TEXT are particular to the database query domain and would have to be redefined if the schemas were to be used in another type of system (such as a tutorial system, for example). The semantics are not particular, however, to the domain of the database. When transferring the system from one database to another, the predicate semantics would not have to be altered.

A proposition is an instantiated predicate; predicate arguments have been filled with values from the knowledge base. An instantiation of the identification predicate is shown below along with its eventual translation.

Instantiated predicate:

(identification (OCEAN-ESCORT CRUISER) SHIP
(non-restrictive TRAVEL-MODE SURFACE))

Eventual translation:

The ocean escort and the cruiser are surface ships.

The schema is filled by stepping through it, using the predicate semantics to select information which matches the predicate arguments. In places where alternative predicates occur in the schema, all alternatives are matched against the relevant knowledge pool producing a set of propositions. The focus constraints are used to select the most appropriate proposition.

The schemas were implemented using a formalism similar to an augmented transition network (ATN). Taking an arc corresponds to the selection of a proposition for the answer. States correspond to filled stages of the schema. The main difference between the TEXT system implementation and a usual ATN, however, is in the control of alternatives. Instead of uncontrolled backtracking, TEXT uses one state lookahead. From a given state, it explores all possible next states and chooses among them using a function that encodes the focus constraints. This use of one state lookahead increases the efficiency of the strategic component since it eliminates unbounded non-determinism.

8.0 FOCUSING MECHANISM

So far, a speaker has been shown to be limited in many ways. For example, s/he is limited by the goal s/he is trying to achieve in the current speech act. TEXT's goal is to answer the user's current question. To achieve that goal, the speaker has limited his/her scope of attention to a set of objects relevant to this goal, as represented by global focus or the relevant knowledge pool. The speaker is also limited by his/her higher-level plan of how to achieve the goal. In TEXT, this plan is the chosen schema. Within these constraints, however, a speaker may still run into the problem of deciding what to say next.

A focusing mechanism is used to provide further constraints on what can be said. The focus constraints used in TEXT are immediate, since they use the most recent proposition (corresponding to a sentence in the English answer) to constrain the next utterance. Thus, as the text is constructed, it is used to constrain what can be said next.

Sidner [SIDNER 79] used three pieces of information for tracking immediate focus: the immediate focus of a sentence (represented by the current focus - CF), the elements of a sentence which are potential candidates for a change in focus (represented by a potential focus list - PFL), and past immediate foci (represented by a focus stack). She showed that a speaker has the following options from one sentence to the next: 1) to continue focusing on the same thing, 2) to focus on one of the items introduced in the last sentence, 3) to return to a previous topic in which case the focus stack is popped, or 4) to focus on an item implicitly related to any of these three options. Sidner's work on focusing concerned the interpretation of anaphora. She says nothing about which of these four options is preferred over others since in interpretation the choice has already been made.

For generation, however, a speaker may have to choose between these options at any point, given all that s/he wants to say. The speaker may be faced with the following choices: 1) continuing to talk about the same thing

(current-focus equals current-focus of the previous sentence) or starting to talk about something introduced in the last sentence (current-focus is a member of potential-focus-list of the previous sentence) and 2) continuing to talk about the same thing (current focus remains the same) or returning to a topic of previous discussion (current focus is a member of the focus-stack).

When faced with the choice of remaining on the same topic or switching to one just introduced, I claim a speaker's preference is to switch. If the speaker has something to say about an item just introduced and does not present it next, s/he must go to the trouble of re-introducing it later on. If s/he does present information about the new item first, however, s/he can easily continue where s/he left off by following Sidner's legal option #3. Thus, for reasons of efficiency, the speaker should shift focus to talk about an item just introduced when s/he has something to say about it.

When faced with the choice of continuing to talk about the same thing or returning to a previous topic of conversation, I claim a speaker's preference is to remain on the same topic. Having at some point shifted focus to the current focus, the speaker has opened a topic for conversation. By shifting back to the earlier focus, the speaker closes this new topic, implying that s/he has nothing more to say about it when in fact, s/he does. Therefore, the speaker should maintain the current focus when possible in order to avoid false implication of a finished topic.

These two guidelines for changing and maintaining focus during the process of generating language provide an ordering on the three basic legal focus moves that Sidner specifies:

1. change focus to member of previous potential focus list if possible -
CF (new sentence) is a member of PFL (last sentence)
2. maintain focus if possible -
CF (new sentence) = CF (last sentence)
3. return to topic of previous discussion -
CF (new sentence) is a member of focus-stack

I have not investigated the problem of incorporating focus moves to items implicitly associated with either current foci, potential focus list members, or previous foci into this scheme. This remains a topic for future research.

Even these guidelines, however, do not appear to be enough to ensure a connected discourse. Although a speaker may decide to focus on a specific entity, s/he may want to convey information about several properties of that entity. S/he will describe related properties of the entity before describing other properties.

Thus, strands of semantic connectivity will occur at more than one level of the discourse.

An example of this phenomenon is given in dialogues (A) and (B) below. In both, the discourse is focusing on a single entity (the balloon), but in (A) properties that must be talked about are presented randomly. In (B), a related set of properties (color) is discussed before the next set (size). (B), as a result, is more connected than (A).

(A) The balloon was red and white striped. Because this balloon was designed to carry men, it had to be large. It had a silver circle at the top to reflect heat. In fact, it was larger than any balloon John had ever seen.

(B) The balloon was red and white striped. It had a silver circle at the top to reflect heat. Because this balloon was designed to carry men, it had to be large. In fact, it was larger than any balloon John had ever seen.

In the generation process, this phenomenon is accounted for by further constraining the choice of what to talk about next to the proposition with the greatest number of links to the potential focus list.

8.1 Use Of The Focus Constraints

TEXT uses the legal focus moves identified by Sidner by only matching schema predicates against propositions which have an argument that can be focused in satisfaction of the legal options. Thus, the matching process itself is constrained by the focus mechanism. The focus preferences developed for generation are used to select between remaining options.

These options occur in TEXT when a predicate matches more than one piece of information in the relevant knowledge pool or when more than one alternative in a schema can be satisfied. In such cases, the focus guidelines are used to select the most appropriate proposition. When options exist, all propositions are selected which have as focused argument a member of the previous PFL. If none exist, then all propositions are selected whose focused argument is the previous current-focus. If none exist, then all propositions are selected whose focused argument is a member of the focus-stack. If these filtering steps do not narrow down the possibilities to a single proposition, that proposition with the greatest number of links to the previous PFL is selected for the answer. The focus and potential focus list of each proposition is maintained and passed to the tactical component for use in selecting syntactic constructions and pronominalization.

Interaction of the focus constraints with the schemas means that although the same schema may be selected for different answers, it can be instantiated in different ways. Recall that the identification schema was selected in response to the question "What is a ship?" and the four predicates, identification, evidence, attributive, and particular-illustration, were instantiated. The identification schema was also selected in response to the question "What is an aircraft carrier?", but different predicates were instantiated as a result of the focus constraints:

(definition AIRCRAFT-CARRIER)

Schema selected: identification

- 1) identification
- 2) analogy
- 3) particular-illustration
- 4) amplification
- 5) evidence

- 1) An aircraft carrier is a surface ship with a DISPLACEMENT between 78000 and 80800 and a LENGTH between 1039 and 1063.
- 2) Aircraft carriers have a greater LENGTH than all other ships and a greater DISPLACEMENT than most other ships.
- 3) Mine warfare ships, for example, have a DISPLACEMENT of 320 and a LENGTH of 144.
- 4) All aircraft carriers in the ONR database have REMARKS of 0, FUEL TYPE of BNKR, FLAG of BLBL, BEAM of 252, ENDURANCE RANGE of 4000, ECONOMIC SPEED of 12, ENDURANCE SPEED of 30 and PROPULSION of SIMTURGRD.
- 5) A ship is classified as an aircraft carrier if the characters 1 through 2 of its HULL_NO are CV.

9.0 FUTURE DIRECTIONS

Several possibilities for further development of the research described here include 1) the use of the same strategies for responding to questions about attributes, events, and relations as well as to questions about entities, 2) investigation of strategies needed for responding to questions about the system processes (e.g. How is manufacturer's cost determined?) or system capabilities (e.g. Can you handle ellipsis?), 3) responding to presuppositional failure as well as to direct questions, and 4) the incorporation of a user model in the generation process (currently TEXT assumes a static casual, naive user and gears its responses to this characterization). This last feature could be used, among other ways, in determining the amount of detail required (see [MCKEOWN 82] for discussion of the recursive use of the schemas).

10.0 CONCLUSION

The TEXT system successfully incorporates principles of relevancy criteria, discourse structure, and focus constraints into a method for generating English text of paragraph length. Previous work on focus of attention has been extended for the task of generation to provide constraints on what to say next. Knowledge about discourse structure has been encoded into schemas that are used to guide the generation process.

The use of these two interacting mechanisms constitutes a departure from earlier generation systems. The approach taken in this research is that the generation process should not simply trace the knowledge representation to produce text. Instead, communicative strategies people are familiar with are used to effectively convey information. This means that the same information may be described in different ways on different occasions.

The result is a system which constructs and orders a message in response to a given question. Although the system was designed to generate answers to questions about database structure (a feature lacking in most natural language database systems), the same techniques and principles could be used in other application areas (for example, computer assisted instruction systems, expert systems, etc.) where generation of language is needed.

Acknowledgements

I would like to thank Aravind Joshi, Bonnie Webber, Kathleen McCoy, and Eric Mays for their invaluable comments on the style and content of this paper. Thanks also goes to Kathleen McCoy and Steven Bossie for their roles in implementing portions of the system.

References

- [BOSSIE 82]. Bossie, S., "A tactical model for text generation: sentence generation using a functional grammar," forthcoming M. S. thesis, University of Pennsylvania, Philadelphia, Pa., 1982.
- [CHEN 76]. Chen, P. P. S., "The entity-relationship model - towards a unified view of data." ACM Transactions on Database Systems, Vol. 1, No. 1. (1976).
- [GRIMES 75]. Grimes, J. E. The Thread of Discourse. Mouton, The Hague, Paris. (1975).
- [GROSZ 77]. Grosz, B. J., "The representation and use of focus in dialogue understanding." Technical note 151, Stanford Research Institute, Menlo Park, Ca. (1977).
- [LEE and GERRITSEN 78]. Lee, R. M. and R. Gerritsen, "Extended semantics for generalization hierarchies", in Proceedings of the 1978 ACM-SIGMOD International Conference on Management of Data, Austin, Tex., 1978.
- [KAY 79]. Kay, M. "Functional grammar." Proceedings of the 5th Annual Meeting of the Berkeley Linguistic Society. (1979).
- [MALHOTRA 75]. Malhotra, A. "Design criteria for a knowledge-based English language system for management: an experimental analysis." MAC TR-146, MIT, Cambridge, Mass. (1975).
- [MCCOY 82]. McCoy, K. F., "Augmenting a database knowledge representation for natural language generation," in Proc. of the 20th Annual Conference of the Association for Computational Linguistics, Toronto, Canada, 1982.
- [MCKEOWN 80]. McKeown, K. R., "Generating relevant explanations: natural language responses to questions about database structure." in Proceedings of AAAI, Stanford Univ., Stanford, Ca. (1980). pp. 306-9.
- [MCKEOWN 82]. McKeown, K. R., "Generating natural language text in response to questions about database structure." Ph.D. dissertation, University of Pennsylvania, Philadelphia, Pa. 1982.
- [SHEPHERD 26]. Shepherd, H. R., The Fine Art of Writing, The Macmillan Co., New York, N. Y., 1926.
- [SIDNER 79]. Sidner, C. L., "Towards a computational theory of definite anaphora comprehension in English discourse." Ph.D. dissertation, MIT AI Technical Report #TR-537, Cambridge, Mass. (1979).
- [SMITH and SMITH 77]. Smith, J. M. and Smith, D. C. P., "Database abstractions: aggregation and generalization." University of Utah, ACM Transactions on Database Systems, Vol. 2, #2, June 1977, pp. 105-33.
- [TENNANT 79]. Tennant, H., "Experience with the evaluation of natural language question answerers." Working paper #18, Univ. of Illinois, Urbana-Champaign, Ill. (1979).