An Information Structural Approach to Spoken Language Generation

Scott Prevost

The Media Laboratory

Massachusetts Institute of Technology

20 Ames Street

Cambridge, Massachusetts 02139-4307 USA

prevost@media.mit.edu

Abstract

This paper presents an architecture for the generation of spoken monologues with contextually appropriate intonation. A two-tiered *information structure* representation is used in the high-level content planning and sentence planning stages of generation to produce efficient, coherent speech that makes certain discourse relationships, such as explicit contrasts, appropriately salient. The system is able to produce appropriate intonational patterns that cannot be generated by other systems which rely solely on word class and given/new distinctions.

1 Introduction

While research on generating coherent *written* text has flourished within the computational linguistics and artificial intelligence communities, research on the generation of *spoken* language, and particularly intonation, has received somewhat less attention. In this paper, we argue that commonly employed models of text organization, such as schemata and rhetorical structure theory (RST), do not adequately address many of the issues involved in generating *spoken* language. Such approaches fail to consider contextually bound focal distinctions that are manifest through a variety of different linguistic and paralinguistic devices, depending on the language.

In order to account for such distinctions of focus, we employ a two-tiered *information structure* representation as a framework for maintaining local coherence in the generation of natural language. The higher tier, which delineates the *theme*, that which links the utterance to prior utterances, and the *rheme*, that which forms the core contribution of the utterance to the discourse, is instrumental in determining the high-level organization of information within a discourse segment. Dividing semantic representations into their thematic and rhematic parts allows propositions to be presented in a way that maximizes the shared material between utterances.

The lower tier in the information structure representation specifies the semantic material that is in “focus” within themes and rhemes. Material may be in focus for a variety of reasons, such as to emphasize its “new” status in the discourse, or to contrast it with other salient material. Such focal distinctions may affect the linguistic presentation of information. For example, the *it*-cleft in (1) may mark *John* as standing in contrast to some other recently mentioned person. Similarly, in (2), the pitch accent on *red* may mark the referenced car as standing in contrast to some other car inferable from the discourse context.¹

(1) It was John who spoke first.

(2) Q: Which car did Mary drive?

A: (MARY drove)_{th} (the RED car.)_{rh}
L+H* LH(%) H* LL\$

By appealing to the notion that the simple rise-fall tune (H* LL%) very often accompanies the rhematic material in an utterance and the rise-fall-rise tune often accompanies the thematic material (Steedman, 1991; Prevost and Steedman, 1994), we present a spoken language generation architecture for producing short spoken monologues with contextually appropriate intonation.

2 Information Structure

Information Structure refers to the organization of information within an utterance. In particular, it

¹In this example, and throughout the remainder of the paper, the intonation contour is informally noted by placing prosodic phrases in parentheses and marking pitch accented words with capital letters. The tunes are more formally annotated with a variant of (Pierrehumbert, 1980) notation described in (Prevost, 1995). Three different pause lengths are associated with boundaries in the modified notation. ‘(%)’ marks intra-utterance boundaries with very little pausing, ‘%’ marks intra-utterance boundaries associated with clauses demarcated by commas, and ‘\$’ marks utterance-final boundaries. For the purposes of generation and synthesis, these distinctions are crucial.

- (3) Q: I know the AMERICAN amplifier produces MUDDY treble,
 (But WHAT) (does the BRITISH amplifier produce?)
 L+H* L(H%) H* LL\$

- (4) Q: I know the AMERICAN amplifier produces MUDDY treble,
 (But WHAT) (produces CLEAN treble?)
 L+H* L(H%) H* LL\$

defines how the information conveyed by a sentence is related to the knowledge of the interlocutors and the structure of their discourse. Sentences conveying the same propositional content in different contexts need not share the same information structure. That is, information structure refers to how the semantic content of an utterance is packaged, and amounts to instructions for updating the models of the discourse participants. The realization of information structure in a sentence, however, differs from language to language. In English, for example, intonation carries much of the burden of information structure, while languages with freer word order, such as Catalan (Engdahl and Vallduví, 1994) and Turkish (Hoffman, 1995) convey information structure syntactically.

2.1 Information Structure and Intonation

The relationship between intonational structure and information structure is illustrated by (3) and (4). In each of these examples, the answer contains the same string words but different intonational patterns and information structural representations. The theme of each utterance is considered to be represented by the material repeated from the question. That is, the theme of the answer is what links it to the question and defines what the utterance is about. The rheme of each utterance is considered to be represented by the material that is new or forms the core contribution of the utterance to the discourse. By mapping the rise-fall tune (H* LL%) onto rhemes and the rise-fall-rise tune (L+H* LH%) onto themes (Steedman, 1991; Prevost and Steedman, 1994), we can easily identify the string of words over which these two prominent tunes occur directly from the information structure. While this mapping is certainly overly simplistic, the results presented in Section 4.3 demonstrate its appropriateness for the class

of simple declarative sentences under investigation.

Knowing the strings of words to which these two tunes are to be assigned, however, does not provide enough information to determine the location of the pitch accents (H* and L+H*) within the tunes. Moreover, the simple mapping described above does not account for the frequently occurring cases in which thematic material bears no pitch accents and is consequently unmarked intonationally. Previous approaches to the problem of determining where to place accents have utilized heuristics based on “givenness.” That is, content-bearing words (e.g. nouns and verbs) which had not been previously mentioned (or whose roots had not been previously mentioned) were assigned accents, while function words were de-accented (Davis and Hirschberg, 1988; Hirschberg, 1990). While these heuristics account for a broad range of intonational possibilities, they fail to account for accentual patterns that serve to contrast entities or propositions that were previously “given” in the discourse. Consider, for example the intonational pattern in (5), in which the pitch accent on *amplifier* in the response cannot be attributed to its being “new” to the discourse.

- (5) Q: Do critics prefer the BRITISH amplifier
 L* H
 or the AMERICAN amplifier?
 H* LL\$
 A: They prefer the AMERICAN amplifier.
 H* LL\$

For the determination of pitch accent placement, we rely on a secondary tier of information structure which identifies focused properties within themes and rhemes. The theme-foci and the rheme-foci mark the information that differentiates properties or entities in the current utterance from properties or entities established in prior utterances. Conse-

(6) Q: Which amplifier does Scott PREFER?
H* LL\$
A: (He prefers)_{th} (the BRITISH amplifier.)_{rh}
H* LL\$

The determination of contrastive focus, and consequently the determination of pitch accent locations, is based on the premise that each object in the knowledge base is associated with a set of alternatives from which it must be distinguished if reference is to succeed. The set of alternatives is determined by the hierarchical structure of the knowledge base. For the present implementation, only properties with the same parent or grandparent class are considered to be alternatives to one another.

(7) *DElist*: a collection of discourse entities that have been evoked in prior discourse, ordered by recency. The list may be limited to some size k so that only the k most recent discourse entities pushed onto the list are retrievable.

$RSet(x, S)$: the set of alternatives for object x as restricted by the referring expressions in $DElist$ and the set of properties S .

Props(*x*): a list of properties for object *x*, ordered by the grammar so that nominal properties take precedence over adjectival properties.

(8) $Props(x) := [P \mid P(x) \text{ is true in KB}]$
 $ASet(x) := \{y \mid alt(x, y)\}$, x 's alternatives
 $RSet(x, \{ \}) := \{x\} \cup \{y \mid y \in ASet(x) \ \& \ y \in DElist\}$, evoked alternatives
 $CSet(x, \{ \}) := \{ \}$

```

(9)  $S := \{ \}$ 
    for each  $P$  in  $Props(x)$ 
         $RSet(x, S \cup \{P\}) :=$ 
             $\{y \mid y \in RSet(x, S) \ \& \ P(y)\}$ 
        if  $RSet(x, S \cup \{P\}) = RSet(x, S)$  then
            % no restrictions were made
            % based on property  $P$ .
             $CSet(x, S \cup \{P\}) := CSet(x, S)$ 
        else
            % property  $P$  eliminated some
            % members of the  $RSet$ .
             $CSet(x, S \cup \{P\}) := CSet(x, S) \cup \{P\}$ 
        endif
     $S := S \cup \{P\}$ 
endfor

```

Based on this contrastive focus algorithm and the mapping between information structure and intonation described above, we can view information structure as the representational bridge between discourse and intonational variability. The following sections elucidate how such a formalism can be integrated into the computational task of generating spoken language.

296

3 Generation Architecture

The task of natural language generation (NLG) has often been divided into three stages: content planning, in which high-level goals are satisfied and discourse structure is determined, sentence planning, in which high-level abstract semantic representations are mapped onto representations that more fully constrain the possible sentential realizations (Rambow and Korelsky, 1992; Reiter and Mellish, 1992; Meteer, 1991), and surface generation, in which the high-level propositions are converted into sentences.

The selection and organization of propositions and their divisions into theme and rheme are determined by the content planner, which maintains discourse coherence by stipulating that semantic information must be shared between consecutive utterances whenever possible. That is, the content planner ensures that the theme of an utterance links it to material in prior utterances.

The process of determining foci within themes and rhemes can be divided into two tasks: determining which discourse *entities* or *propositions* are in focus, and determining how their linguistic realizations should be marked to convey that focus. The first of these tasks can be handled in the content phase of the NLG model described above. The second of these tasks, however, relies on information, such as the construction of referring expressions, that is often considered the domain of the sentence planning stage. For example, although two discourse entities e_1 and e_2 can be determined to stand in contrast to one another by appealing only to the discourse model and the salient pool of knowledge, the method of contrastively distinguishing between them by the placement of pitch accents cannot be resolved until the choice of referring expressions has been made. Since referring expressions are generally taken to be in the domain of the sentence planner (Dale and Haddock, 1991), the present approach resolves issues of contrastive focus assignment at the sentence processing stage as well.

During the content generation phase, the content of the utterance is planned based on the previous discourse. While template-based systems (McKeown, 1985) have been widely used, rhetorical structure theory (RST) approaches (Mann and Thompson, 1986; Hovy, 1993), which organize texts by identifying rhetorical relations between clause-level propositions from a knowledge base, have recently flourished. Sibun (Sibun, 1991) offers yet another alternative in which propositions are linked to one another not by rhetorical relations or pre-planned templates, but rather by physical and spatial properties represented in the knowledge-base.

The present framework for organizing the content of a monologue is a hybrid of the template and RST approaches. The implementation, which is presented in the following section, produces de-

scriptions of objects from a knowledge base with context-appropriate intonation that makes proper distinctions of contrast between alternative, salient discourse entities. Certain constraints, such as the requirement that objects be identified or defined at the beginning of a description, are reminiscent of McKeown's schemata. Rather than imposing strict rules on the order in which information is presented, the order is determined by domain specific knowledge, the communicative intentions of the speaker, and beliefs about the hearer's knowledge. Finally, the system includes a set of rhetorical constraints that may rearrange the order of presentation for information in order to make certain rhetorical relationships salient. While this approach has proven effective in the present implementation, further research is required to determine its usefulness for a broader range of discourse types.

4 The Prolog Implementation

The monologue generation program produces text and contextually-appropriate intonation contours to describe an object from the knowledge base. The system exhibits the ability to intonationally contrast alternative entities and properties that have been explicitly evoked in the discourse even when they occur with several intervening sentences.

4.1 Content Generation

The architecture for the monologue generation program is shown in Figure 1, in which arrows represent the computational flow and lines represent dependencies among modules. The remainder of this section contains a description of the computational path through the system with respect to a single example. The input to the program is a goal to describe an object from the knowledge base, which in this case contains a variety of facts about hypothetical stereo components. In addition, the input provides a communicative intention for the goal which may affect its ultimate realization, as shown in (10). For example, given the goal `describe(x)`, the intention `persuade-to-buy(hearer,x)` may result in a radically different monologue than the intention `persuade-to-sell(hearer,x)`.

- (10) *Goal: describe e1*
Input: generate(intention(bel(h1,
good-to-buy(e1)))

Information from the knowledge base is selected to be included in the output by a set of relations that determines the degree to which knowledge base facts and rules support the communicative intention of the speaker. For example, suppose the system "believes" that conveying the proposition in (11) moderately supports the intention of making hearer *h1* want to buy *e1*, and further that the rule in (12) is known by *h1*.

Figure 1: An Architecture for Monologue Generation

- ```

(11) bel(h1, holds(rating(X, powerful)))
(12) holds(rating(X, powerful)) :-
 holds(produce(X, Y)),
 holds(isa(Y, watts-per-channel)),
 holds(amount(Y, Z)),
 number(Z),
 Z >= 100.

```

The program then consults the facts in the knowledge base, verifies that the property does indeed hold and consequently includes the corresponding facts in the set of properties to be conveyed to the hearer, as shown in (13).

- ```

(13) holds(produce(e1, e7)).
    holds(isa(e7, watts-per-channel)).
    holds(amount(e7, 100)).
  
```

The content generator starts with a simple description template that specifies that an object is to be explicitly identified or defined before other propositions concerning it are put forth. Other relevant propositions concerning the object in question are then linearly organized according to beliefs about how well they contribute to the overall intention. Finally, a small set of rhetorical predicates rearranges the linear ordering of propositions so that sets of sentences that stand in some interesting rhetorical relationship to one another will be realized together in the output. These rhetorical predicates employ information structure to assist in maintaining the coherence of the output. For example, the *conjunction* predicate specifies that propositions sharing the same theme or rheme be realized together in order to avoid excessive topic shifting. The *contrast* predicate specifies that pairs of themes or rhemes that explicitly contrast with one another be realized to-

gether. The result is a set of properties roughly ordered by the degree to which they support the given intention, as shown in (14).

- ```

(14) holds(defn(isa(e1, amplifier)))
 holds(design(e1, solid-state), pres)
 holds(cost(e1, e9), pres)
 holds(produce(e1, e7), pres)
 holds(contrast(praise(e4, e1),
 revile(e5, e1)), past)

```

The top-level propositions shown in (14) were selected by the program because the hearer (h1) is believed to be interested in the design of the amplifier and the reviews the amplifier has received. Moreover, the belief that the hearer is interested in buying an expensive, powerful amplifier justifies including information about its cost and power rating. Different sets of propositions would be generated for other (perhaps thriftier) hearers. Additionally, note that the propositions *praise(e4, e1)* and *revile(e5, e1)* are combined into the larger proposition *contrast(praise(e4, e1), revile(e5, e1))*. This is accomplished by the rhetorical constraints that determine the two propositions to be contrastive because *e4* and *e5* belong to the same set of alternative *entities* in the knowledge base and *praise* and *revile* belong to the same set of alternative *propositions* in the knowledge base.

The next phase of content generation recognizes the dependency relationships between the properties to be conveyed based on shared discourse entities. This phase, which represents an extension of the rhetorical constraints, arranges propositions to ensure that consecutive utterances share semantic material (cf. (McKeown et al., 1994)). This rule, which in effect imposes a strong bias for Centering Theory's *continue* and *retain* transitions (Grosz et al., 1986) determines the theme-rheme segmentation for each proposition.

## 4.2 Sentence Planning

After the coherence constraints from the previous section are applied, the sentence planner is responsible for making decisions concerning the form in which propositions are realized. This is accomplished by the following simple set of rules. First, Definitional *isa* properties are realized by the matrix verb. Other *isa* properties are realized by nouns or noun phrases. Top-level properties (such as those in (14)) are realized by the matrix verb. Finally, embedded properties (those evoked for building referring expressions for discourse entities) are realized by adjectival modifiers if possible and otherwise by relative clauses.

While there are certainly a number of linguistically interesting aspects to the sentence planner, the most important aspect for the present purposes is the determination of theme-foci and rheme-foci. The focus assignment algorithm employed by the

sentence planner, which has access to both the discourse model and the knowledge base, works as follows. First, each property or discourse entity in the semantic and information structural representations is marked as either previously mentioned or new to the discourse. This assignment is made with respect to two data structures, the discourse entity list (DEList), which tracks the succession of entities through the discourse, and a similar structure for evoked properties. Certain aspects of the semantic form are considered unaccentable because they correspond to the interpretations of closed-class items such as function words. Items that are assigned focus based on their “newness” are assigned the  $\circ$  focus operator, as shown in (15).

- (15) Semantics:  $defn(isa(\circ e1, \circ c1))$   
 Theme:  $\circ e1$   
 Rheme:  $\lambda x. isa(x, \circ c1)$   
 Supporting Props:  $isa(c1, \circ amplifier)$   
 $\circ design(c1, \circ solidstate)$

The second step in the focus assignment algorithm checks for the presence of contrasting propositions in the ISStore, a structure that stores a history of information structure representations. Propositions are considered contrastive if they contain two contrasting pairs of discourse entities, or if they contain one contrasting pair of discourse entities as well as contrasting functors.

Discourse entities are determined to be contrastive if they belong to the same set of alternatives in the knowledge base, where such sets are inferred from the *isa*-links that define class hierarchies. While the present implementation only considers entities with the same parent or grandparent class to be alternatives for the purposes of contrastive stress, a graduated approach that entails degrees of contrastiveness may also be possible.

The effects of the focus assignment algorithm are easily shown by examining the generation of an utterance that contrasts with the utterance shown in (15). That is, suppose the generation program has finished generating the output corresponding to the examples in (10) through (15) and is assigned the new goal of describing entity *e2*, a different amplifier. After applying the second step on the focus assignment algorithm, contrasting discourse entities are marked with the  $\bullet$  contrastive focus operator, as shown in (16). Since *e1* and *e2* are both instances of the class *amplifiers* and *c1* and *c2* both describe the class *amplifiers* itself, these two pairs of discourse entities are considered to stand in contrastive relationships.

- (16) Semantics:  $defn(isa(\bullet e2, \bullet c2))$   
 Theme:  $\bullet e2$   
 Rheme:  $\lambda x. isa(x, \bullet c2)$   
 Supporting Props:  $class(c2, amplifier)$   
 $design(c2, \circ tube)$

While the previous step of the algorithm determined which abstract discourse entities and properties stand in contrast, the third step uses the contrastive focus algorithm described in Section 2 to determine which elements need to be contrastively focused for reference to succeed. This algorithm determines the minimal set of properties of an entity that must be “focused” in order to distinguish it from other salient entities. For example, although the representation in (16) specifies that *e2* stands in contrast to some other entity, it is the property of *e2* having a tube design rather than a solid-state design that needs to be conveyed to the hearer. After applying the third step of the focus assignment to (16), the result appears as shown in (17), with “tube” contrastively focused as desired.

- (17) Semantics:  $defn(isa(\bullet e2, \bullet c2))$   
 Theme:  $\bullet e2$   
 Rheme:  $\lambda x. isa(x, \bullet c2)$   
 Supporting Props:  $isa(c2, amplifier)$   
 $design(c2, \bullet tube)$

The final step in the sentence planning phase of generation is to compute a representation that can serve as input to a surface form generator based on Combinatory Categorical Grammar (CCG) (Steedman, 1991), as shown in (18).<sup>3</sup>

- (18) Theme:  $np(3, s) :$   
 $(e1 \wedge S) \wedge def(e1, \bullet x5(e1) \& S) @ u/rh$   
 Rheme:  $s :$   
 $(act \wedge pres) \wedge indef(c1, (amplifier(c1) \& \bullet tube(c1)) \& isa(e1, c1)) \backslash np(3, s) : e1 @ rh$

### 4.3 Results

Given the focus-marked output of the sentence planner, the surface generation module consults a CCG grammar which encodes the information structure/intonation mapping and dictates the generation of both the syntactic and prosodic constituents. The result is a string of words and the appropriate prosodic annotations, as shown in (19). The output of this module is easily translated into a form suitable for a speech synthesizer, which produces spoken output with the desired intonation.<sup>4</sup>

- (19) The X5 is a TUBE amplifier.  
 $L+H^* \quad L(H\%) \quad H^* \quad LL\$$

The modules described above and shown in Figure 1 are implemented in Quintus Prolog. The system produces the types of output shown in (20) and

<sup>3</sup>A complete description of the CCG generator can be found in (Prevost and Steedman, 1993). CCG was chosen as the grammatical formalism because it licenses non-traditional syntactic constituents that are congruent with the bracketings imposed by information structure and intonational phrasing, as illustrated in (3).

<sup>4</sup>The system currently uses the AT&T Bell Laboratories TTS system, but the implementation is easily adaptable to other synthesizers.

(21), which should be interpreted as a single (two paragraph) monologue satisfying a goal to describe two different objects.<sup>5</sup> Note that both paragraphs include very similar types of information, but radically different intonational contours, due to the discourse context. In fact, if the intonational patterns of the two examples are interchanged, the resulting speech sounds highly unnatural.

- (20) a. Describe the x4.  
 b. The X4  
     L+H\* L(H%)  
 is a SOLID-state AMPLIFIER.  
     H\*           H\*           LL\$  
 It COSTS EIGHT HUNDRED DOLLARS,  
     H\*       H\*       H\*           H\* LL%  
 and PRODUCES  
     H\*  
 ONE hundred watts-per-CHANNEL.  
     H\*                           H\* LL\$  
 It was PRAISED by STEREOFOL,  
     H<sub>c</sub>           !H<sub>c</sub>           LH%  
 an AUDIO JOURNAL,  
     H\*       H\*       LH%  
 but was REVILED by AUDIOFAD,  
     H<sub>c</sub>           !H<sub>c</sub>           LH%  
 ANOTHER audio journal.  
     H\*                           LL\$

- (21) a. Describe the x5.  
 b. The X5                   is a TUBE amplifier.  
     L+H<sub>c</sub> L(H%)           H<sub>c</sub>           LL\$  
 IT           costs NINE hundred dollars,  
     L+H<sub>c</sub> L(H%)   H<sub>c</sub>           LH%  
 produces TWO hundred watts-per-channel.  
     H<sub>c</sub>                           LH%  
 and was praised  
 by Stereofool AND Audiofad.  
     H<sub>c</sub>                           LL\$

Several aspects of the output shown above are worth noting. Initially, the program assumes that the hearer has no specific knowledge of any particular objects in the knowledge base. Note however, that every proposition put forth by the generator is assumed to be incorporated into the hearer's set of beliefs. Consequently, the descriptive phrase "an audio journal," which is new information in the first paragraph, is omitted from the second. Additionally, when presenting the proposition '*Audiofad is an audio journal*,' the generator is able to recognize the similarity with the corresponding proposition about Stereofool (i.e. both propositions are abstractions over the single variable open proposition '*X is an audio journal*'). The program therefore interjects the *other* property and produces "another audio journal."

<sup>5</sup>The implementation assigns slightly higher pitch to accents bearing the subscript *c* (e.g. H<sub>c</sub>\*), which mark contrastive focus as determined by the algorithm describe above and in (Prevost, 1995).

Several aspects of the contrastive intonational effects in these examples also deserve attention. Because of the content generator's use of the rhetorical contrast predicate, items are eligible to receive stress in order to convey contrast before the contrasting items are even mentioned. This phenomenon is clearly illustrated by the clause "PRAISED by STEREOFOL" in (20), which is contrastively stressed before "REVILED by AUDIOFAD" is uttered. Such situations are produced only when the contrasting propositions are gathered by the content planner in a single invocation of the generator and identified as contrastive when the rhetorical predicates are applied. Moreover, unlike systems that rely solely on word class and given/new distinctions for determining accentual patterns, the system is able to produce contrastive accents on pronouns despite their "given" status, as shown in (21).

## 5 Conclusions

The generation architecture described above and implemented in Quintus Prolog produces paragraph-length, spoken monologues concerning objects in a simple knowledge base. The architecture relies on a mapping between a two-tiered information structure representation and intonational tunes to produce speech that makes appropriate contrastive distinctions prosodically. The process of natural language generation, in accordance with much of the recent literature in the field, is divided into three processes: high-level content planning, sentence planning, and surface generation. Two points concerning the role of intonation in the generation process are emphasized. First, since intonational phrasing is dependent on the division of utterances into theme and rheme, and since this division relates consecutive sentences to one another, matters of information structure (and hence intonational phrasing) must be largely resolved during the high-level planning phase. Second, since accentual decisions are made with respect to the particular linguistic realizations of discourse properties and entities (e.g. the choice of referring expressions), these matters cannot be fully resolved until the sentence planning phase.

## 6 Acknowledgments

The author is grateful for the advice and helpful suggestions of Mark Steedman, Justine Cassell, Kathy McKeown, Aravind Joshi, Ellen Prince, Mark Liberman, Matthew Stone, Beryl Hoffman and Kris Thórisson as well as the anonymous ACL reviewers. Without the AT&T Bell Laboratories TTS system, and the patient advice on its use from Julia Hirschberg and Richard Sproat, this work would not have been possible. This research was funded by NSF grants IRI91-17110 and IRI95-04372 and the generous sponsors of the MIT Media Laboratory.

## References

- Bolinger, D. (1989). *Intonation and Its Uses*. Stanford University Press.
- Culicover, P. and Rochemont, M. (1983). Stress and focus in English. *Language*, 59:123–165.
- Dale, R. and Haddock, N. (1991). Content determination in the generation of referring expressions. *Computational Intelligence*, 7(4):252–265.
- Davis, J. and Hirschberg, J. (1988). Assigning intonational features in synthesized spoken discourse. In *Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics*, pages 187–193, Buffalo.
- Engdahl, E. and Vallduví, E. (1994). Information packaging and grammar architecture: A constraint-based approach. In Engdahl, E., editor, *Integrating Information Structure into Constraint-Based and Categorical Approaches (DYANA-2 Report R.1.3.B)*. CLLI, Amsterdam.
- Grosz, B. J., Joshi, A. K., and Weinstein, S. (1986). Towards a computational theory of discourse interpretation. Unpublished manuscript.
- Gussenhoven, C. (1983b). *On the Grammar and Semantics of Sentence Accent*. Foris, Dordrecht.
- Halliday, M. (1970). Language structure and language function. In Lyons, J., editor, *New Horizons in Linguistics*, pages 140–165. Penguin.
- Hirschberg, J. (1990). Accent and discourse context: Assigning pitch accent in synthetic speech. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, pages 952–957.
- Hoffman, B. (1995). *The Computational Analysis of the Syntax and Interpretation of 'Free' Word Order in Turkish*. PhD thesis, University of Pennsylvania, Philadelphia.
- Hovy, E. (1993). Automated discourse generation using discourse structure relations. *Artificial Intelligence*, 63:341–385.
- Mann, W. and Thompson, S. (1986). Rhetorical structure theory: Description and construction of text structures. In Kempen, G., editor, *Natural Language Generation: New Results in Artificial Intelligence, Psychology and Linguistics*, pages 279–300. Kluwer Academic Publishers, Boston.
- McKeown, K., Kukich, K., and Shaw, J. (1994). Practical issues in automatic documentation generation. In *Proceedings of the Fourth ACL Conference on Applied Natural Language Processing*, pages 7–14, Stuttgart. Association for Computational Linguistics.
- McKeown, K. R. (1985). *Text Generation: Using Discourse Strategies and Focus Constraints to Generate Natural Language Text*. Cambridge University Press, Cambridge.
- Meteer, M. (1991). Bridging the generation gap between text planning and linguistic realization. *Computational Intelligence*, 7(4):296–304.
- Pierrehumbert, J. (1980). *The Phonology and Phonetics of English Intonation*. PhD thesis, Massachusetts Institute of Technology. Distributed by Indiana University Linguistics Club, Bloomington, IN.
- Prevost, S. (1995). *A Semantics of Contrast and Information Structure for Specifying Intonation in Spoken Language Generation*. PhD Thesis, University of Pennsylvania.
- Prevost, S. and Steedman, M. (1993). Generating contextually appropriate intonation. In *Proceedings of the 6th Conference of the European Chapter of the Association for Computational Linguistics*, pages 332–340, Utrecht.
- Prevost, S. and Steedman, M. (1994). Specifying intonation from context for speech synthesis. *Speech Communication*, 15:139–153.
- Rambow, O. and Korelsky, T. (1992). Applied text generation. In *Proceedings of the Third Conference on Applied Natural Language Processing (ANLP-1992)*, pages 40–47.
- Reiter, E. and Mellish, C. (1992). Using classification to generate text. In *Proceedings of the 30th Annual Meeting of the Association for Computational Linguistics*, pages 265–272.
- Robin, J. (1993). A revision-based generation architecture for reporting facts in their historical context. In Horacek, H. and Zock, M., editors, *New Concepts in Natural Language Generation: Planning, Realization and Systems*, pages 238–265. Pinter Publishers, New York.
- Rochemont, M. (1986). *Focus in Generative Grammar*. John Benjamins, Philadelphia.
- Sibun, P. (1991). *The Local Organization and Incremental Generation of Text*. PhD thesis, University of Massachusetts.
- Steedman, M. (1991a). Structure and intonation. *Language*, pages 260–296.