A Constraint-based Approach to English Prosodic Constituents

Ewan Klein

Division of Informatics
University of Edinburgh
2 Buccleuch Place, Edinburgh EH8 9LW, UK
ewan@cogsci.ed.ac.uk

Abstract

The paper develops a constraint-based theory of prosodic phrasing and prominence, based on an HPSG framework, with an implementation in ALE. Prominence and juncture are represented by n -ary branching metrical trees. The general aim is to define prosodic structures recursively, in parallel with the definition of syntactic structures. We address a number of *prima facie* problems arising from the discrepancy between syntactic and prosodic structure

1 Introduction

This paper develops a declarative treatment of prosodic constituents within the framework of constraint-based phonology, as developed for example in (Bird, 1995; Mastroianni and Carpenter, 1994). On such an approach, phonological representations are encoded with typed feature terms. In addition to the representational power of complex feature values, the inheritance hierarchy of types provides a flexible mechanism for classifying linguistic structures, and for expressing generalizations by means of type inference.

To date, little work within constraint-based phonology has addressed prosodic structure above the level of the foot. In my treatment, I will adopt the following assumptions:

1. Phonology is induced in parallel with syntactic structure, rather than being mapped from pre-built parse trees.
2. Individual lexical items do not impose constraints on their neighbour's phonology.

The first of these assumptions ensures that phonology is compositional, in the sense that the only information available when assembling the phonology of a complex constituent is the phonology of that constituent's daughters. The second assumption is one that is standardly adopted in HPSG (Pollard and Sag, 1994),

in the sense that heads can be subcategorized with respect to the syntactic and semantic properties of their arguments (i.e., their arguments' *synsem* values), but not with respect to their arguments' phonological properties. Although I am not convinced that this restriction is correct, it is worthwhile to explore what kinds of phonological analyses are compatible with it.

Most of the data used in this paper was drawn from the SOLE spoken corpus (Hitzeman et al., 1998).¹ The corpus was based on recordings of one speaker reading approximately 40 short descriptive texts concerning jewelry.

2 Syntactic and Prosodic Structure

2.1 Metrical Trees

Metrical trees were introduced by Liberman (1977) as a basis for formulating stress-assignment rules in both words and phrases. Syntactic constituents are assigned relative prosodic weight according to the following rule:

- (1) *NSR*: In a configuration [_C A B], if *C* is a phrasal category, *B* is strong.

Prominence is taken to be a relational notion: a constituent labelled 's' is stronger than its sister. Consequently, if *B* in (1) is strong, then *A* must be weak.

In the case of a tree like (2), Liberman and Prince's (1) yields a binary-branching structure of the kind illustrated in (3) (where the root of the tree is unlabeled):

¹The task of recovering relevant examples from the SOLE corpus was considerably aided by the Gsearch corpus query system (Corley et al., 1999).

For any given constituent analysed by a metrical tree t , the location of its main stress can be found by tracing a path from the root of t to a terminal element α such that all nodes on that path are labelled 's'. Thus the main stress in (3) is located on the element *cloak*. In general, the most prominent element, defined in this way, is called the *Designated Terminal Element* (DTE) (Lieberman and Prince, 1977).

Note that (1) is the metrical version of Chomsky and Halle's (1968) Nuclear Stress Rule (NSR), and encodes the same claim, namely that in the default case, main stress falls on the last constituent in a given phrase. Of course, it has often been argued that the notion of 'default prominence' is flawed, since it supposes that the acceptability of utterances can be judged in a null context. Nevertheless, there is an alternative conception: the predictions of the NSR correctly describe the prominence patterns when the whole proposition expressed by the clause in question receives broad focus (Ladd, 1996). This is the view that I will adopt. Although I will concentrate in the rest of the paper on the broad focus pattern of intonation, the approach I develop is intended to link up eventually with pragmatic information about the location of narrow focus.

In the formulation above, (1) only applies to binary-branching constituents, and the question arises how non-binary branching constituent structures (e.g., for VPs headed by ditransitive verbs) should be treated. One option (Beckman, 1986; Pierrehumbert and Beckman, 1988; Nespor and Vogel, 1986) would be to drop the restriction that metrical trees are binary, allowing structures such as Fig 1. Since the nested structure which results from binary branching appears to be irrelevant to phonetic interpretation, I will use n -ary metrical trees in the following analysis.

In this paper, I will not make use of the Prosodic Hierarchy (Beckman and Pierrehumbert, 1986; Nespor and Vogel, 1986; Selkirk, 1981; Selkirk, 1984). Most of the phenomena that I wish to deal with lie in the blurry region (Shattuck-Hufnagel and Turk, 1996) between the Phonological Word and the Intonational Phrase (IP), and I will just refer to 'prosodic constituents' without committing myself to a specific set of labels. I will also not adopt the Strict Layer Hypothesis (Selkirk, 1984) which holds that elements of a given prosodic category (such as Intonational Phrase) must be exhaustively

analysed into a sequence of elements of the next lower category (such as Phonological Phrase). However, it is important to note that every IP will be a prosodic constituent, in my sense. Moreover, my lower-level prosodic constituents could be identified with the ϕ -phrases of (Selkirk, 1981; Gee and Grosjean, 1983; Nespor and Vogel, 1986; Bachenko and Fitzpatrick, 1990), which are grouped together to make IPs.

2.2 Representing Prosodic Structure

I shall follow standard assumptions in HPSG by separating the phonology attribute out from syntax-semantics (SYNSEM):

$$(4) \text{ feat-struct} \rightarrow \begin{bmatrix} \text{PHON} & \text{pros} \\ \text{SYNSEM} & \text{synsem} \end{bmatrix}$$

The type of value of PHON is *pros* (i.e., prosody). In this paper, I am going to take word forms as phonologically simple. This means that the prosodic type of word forms will be maximal in the hierarchy. The only complex prosodic objects will be metrical trees. The minimum requirements for these are that we have, first, a way of representing nested prosodic domains, and second, a way of marking the strong element (Designated Terminal Element; DTE) in a given domain.

Before elaborating the prosodic signature further, I need to briefly address the prosodic status of monosyllabic function words in English. Although these are sometimes classified as clitics, Zwicky (1982) proposes the term *Leaners*. These "form a rhythmic unit with the neighbouring material, are normally unstressed with respect to this material, and do not bear the intonational peak of the unit. English articles, coordinating conjunctions, complementizers, relative markers, and subject and object pronouns are all leaners in this sense" (Zwicky, 1982, p5). Zwicky takes pains to differentiate between Leaners and clitics; the former combine with neighbours to form Phonological Phrases (with juncture characterized by external sandhi), whereas clitics combine with their hosts to form Phonological Words (where juncture is characterized by internal sandhi).

Since Leaners cannot bear intonational peaks, they cannot act as the DTE of a metrical tree. Consequently, the value of the attribute DTE in a metrical tree must be the type of all prosodic objects which are not Leaners. I call this type *full*, and it subsumes both Prosodic Words (of type *p-wrd*) and metrical trees (of type *mtr*). Moreover, since Leaners form a closer juncture with their neighbours than Prosodic Words do, we distinguish two kinds of metrical tree. In a tree of type *full-mtr*, all the daughters are of type *full*, whereas in a tree of type *lnr-mtr*, only the DTE is of type *full*.

$$\left[\begin{array}{l} \text{phrase} \\ \text{PHON} \quad \text{mkMtr}(\langle \phi_0, \phi_1, \dots, \phi_n \rangle) \\ \text{SYNSEM} \quad \left[\begin{array}{l} \text{COMPS} \quad \langle \rangle \end{array} \right] \end{array} \right] \rightarrow \left[\begin{array}{l} \text{word} \\ \text{PHON} \quad \phi_0 \\ \text{COMPS} \quad \left\langle \begin{array}{l} \boxed{1}[\text{PHON } \phi_1], \dots, \boxed{n}[\text{PHON } \phi_n] \end{array} \right\rangle \end{array} \right] \boxed{1}, \dots, \boxed{n}$$

Figure 4: Head-Complement Rule

Complement Rule (Pollard and Sag, 1994), illustrated in Fig 4. The ALE rendering of the rule is given in (7).

```
(7) head_complement_rule
    (phrase, phon:MoPhon,
     synsem:(comps:[],
              spr:S,
              head:Head))
    ==>
    cat> (word, phon:HdPhon,
          synsem:(comps:Comps,
                  spr:S,
                  head:Head)),
    cats> Comps,
    goal> (getPhon(Comps, PhonList),
           mkMtr([HdPhon|PhonList], MoPhon)).
```

The function *mkMtr* (*make metrical tree*) (encoded as a relational constraint in (7)) takes a list consisting of all the daughters' phonologies and builds an appropriate prosodic object ϕ . As the name of the function suggests, this prosodic object is, in the general case, a metrical tree. However, since metrical trees are relational (i.e., one node is *stronger* than the others), it makes no sense to construct a metrical tree if there is only a single daughter. In other words, if the head's COMPS list is empty, then the argument *mkMtr* is a singleton list containing only the head's PHON value, and this is returned unaltered as the function value.

(8) $\text{mkMtr}(\langle \boxed{1}[\text{pros}] \rangle) = \boxed{1}$

The general case requires at least the first two elements on the list of prosodies to be of type *full*, and builds a tree of type *full_mtr*.

(9) $\text{mkMtr}(\langle \boxed{1}[\text{full}], [\text{full}], \dots, \boxed{2} \rangle) =$

$$\left[\begin{array}{l} \text{full-mtr} \\ \text{DOM} \quad \boxed{1} \\ \text{DTE} \quad \boxed{2} \end{array} \right]$$

Note that the domain of the output tree is the input list, and the DTE is just the right-hand element of the domain. (10) shows the constraint in ALE notation; the relation *rhD_DTE/2* simply picks out the last element of the list *L*.

```
(10) mkMtr((full, full|_, L),
           (full_mtr, dom:L, dte:X)) if
           rhD_DTE(L, X).
```

Examples of the prosody constructed for an N-bar and a VP are illustrated in (11)–(12). For convenience, I use [*of the samurai*] to abbreviate the AVM representation of the metrical tree for *of the samurai*, and similarly for [*a cloak*] and [*at the collar*].

(11) $\text{mkMtr}(\langle \text{possession}, [\text{of the samurai}] \rangle) =$

$$\left[\begin{array}{l} \text{full-mtr} \\ \text{DOM} \quad \langle \text{possession}, \boxed{1}[\text{of the samurai}] \rangle \\ \text{DTE} \quad \boxed{1} \end{array} \right]$$

(12) $\text{mkMtr}(\langle \text{fasten}, [\text{a cloak}], [\text{at the collar}] \rangle) =$

$$\left[\begin{array}{l} \text{full-mtr} \\ \text{DOM} \quad \langle \text{fasten}, [\text{a cloak}], \boxed{1}[\text{at the collar}] \rangle \\ \text{DTE} \quad \boxed{1} \end{array} \right]$$

Let's now briefly consider the case of a weak pronominal NP occurring within a VP. Zwicky (1986) develops a prosodically-based account of the distribution of unaccented pronouns in English, as illustrated in the following contrasts:

- (13) a. We took in the unhappy little mutt right away.
 b. *We took in *him* right away.
 c. We took *him* in right away.
- (14) a. Martha told Noel the plot of *Gravity's Rainbow*.
 b. *Martha told Noel *it*.
 c. Martha told *it* to Noel.

Pronominal NPs can only form prosodic phrases in their own right if they bear accent; unaccented pronominals must combine with a host to be admissible. Zwicky's constraints on when this combination can occur are as follows:

- (15) A personal pronoun NP can form a prosodic phrase with a preceding prosodic host only if the following conditions are satisfied:
- the prosodic host and the pronominal NP are sisters;
 - the prosodic host is a lexical category;
 - the prosodic host is a category that governs case marking.

$$\begin{bmatrix} \text{phrase} \\ \text{PHON} & \text{extMtr}(\varphi_1, \varphi_0) \\ \text{SYNSEM} & \begin{bmatrix} \text{SPR} & \langle \rangle \end{bmatrix} \end{bmatrix} \rightarrow \boxed{1} \begin{bmatrix} \text{phrase} \\ \text{PHON} & \varphi_0 \\ \text{SPR} & \langle \boxed{1} [\text{PHON } \varphi_1] \rangle \end{bmatrix}$$

Figure 5: Head-Specifier Rule

These considerations motivate a third clause to the definition of *mkMtr*:

$$(16) \text{mkMtr}(\langle \boxed{1}[p\text{-}wrd], \boxed{2}[lnr] \rangle \oplus \boxed{3}) = \text{mkMtr}(\langle \begin{bmatrix} lnr\text{-}mtr \\ \text{DOM} & \langle \boxed{1}, \boxed{2} \rangle \\ \text{DTE} & \boxed{1} \end{bmatrix} \oplus \boxed{3} \rangle)$$

That is, if the first two elements of the list are a Prosodic Word and a Leaner, then the two of them combine to form a *lnr-mtr*, followed by any other material on the input list. Because of the way in which this prosodic constraint is associated with the Head-Complement Rule, the prosodic host in (16), namely the *p-wrd* tagged $\boxed{1}$, is automatically the syntactic head of the construction. As a result, Zwicky's conditions in (15) fall out directly.

(17)–(18) illustrate the effects of the new clause. In the first case, the *lnr-mtr* consisting of *told* and *it* is the only item on the list in the recursive call to *mkMtr* in (16), and hence the base clause (8) in the definition of *mkMtr* applies. In the second case, there is more than one item on the list, and the *lnr-mtr* becomes a subtree in a larger metrical domain.

$$(17) \text{mkMtr}([told, it]) = \begin{bmatrix} lnr\text{-}mtr \\ \text{DOM} & \langle \boxed{1}told, it \rangle \\ \text{DTE} & \boxed{1} \end{bmatrix}$$

$$(18) \text{mkMtr}([told, it, [to Noel]]) = \begin{bmatrix} full\text{-}mtr \\ \text{DOM} & \langle \begin{bmatrix} lnr\text{-}mtr \\ \text{DOM} & \langle \boxed{1}told, it \rangle \\ \text{DTE} & \boxed{1} \end{bmatrix} \boxed{2}[to Noel] \rangle \\ \text{DTE} & \boxed{2} \end{bmatrix}$$

By contrast, examples of the form *told Noel* *it* fail to parse, since (16) only licenses a head-initial *lnr-mtr* when the Leaner immediately follows the head. We could however admit *told Noel* *it*, if the lexicon contained a suitable entry for accent-bearing *it* with prosody of type *p-wrd*, since this would satisfy the requirement that only prosodies of type *full* can be the value of a metrical tree's DTE.

3.2 Head-Final Constructions

To illustrate head-final constructions, I will focus on NP structures, considering the combination of determiners and pronominal adjectives with N-bar phrases. I take the general case to be illustrated by combining a determiner like *this* with a phrase like *treasured possession* to form one metrical tree. Since *treasured possession* will itself be a metrical tree, I introduce a new, binary, function for this purpose, namely *extMtr* (*extend* metrical tree) which adds a new prosodic element to the left boundary of an existing tree. For convenience, I will call the leftmost argument of *extMtr* the *extender*.

Fig 5 illustrates the way in which *extMtr* is used to build the prosody of a specifier-head construction, while (19) provides the definition of *extMtr*. An example of the output is illustrated in (20).

$$(19) \text{extMtr}(\boxed{1}[full], \begin{bmatrix} \text{DOM} & \boxed{2} \\ \text{DTE} & \boxed{3} \end{bmatrix}) = \begin{bmatrix} full\text{-}mtr \\ \text{DOM} & \boxed{1} \oplus \boxed{2} \\ \text{DTE} & \boxed{3} \end{bmatrix}$$

$$(20) \text{extMtr}(this, [treasured possession]) = \begin{bmatrix} full\text{-}mtr \\ \text{DOM} & \langle this, treasured, \boxed{1}possession \rangle \\ \text{DTE} & \boxed{1} \end{bmatrix}$$

However, there are now a number of special cases to be considered. First, we have to allow that the head phrase is a single Prosodic Word such as *possession*, rather than a metrical tree. Second, the prosodic structure to be built will be more complex if the head phrase itself contains a post-head complement, as in *treasured possession of the samurai*. Crosscutting this dimension is the question of whether the extender is a Leaner, in which case it will form a *lnr-mtr* with the immediately following element. We will look at these cases in turn.

(i) The head is a single Prosodic Word When the second prosodic argument of *extMtr* is not in fact a metrical tree, it calls *mkMtr* to build a new metrical tree. Definition (21) is illustrated in (22).

Figure 6: Right-branching NP Structure

Figure 7: Flat NP Prosodic Structure

$$(21) \text{extMtr}(\boxed{1}[\text{pros}], \boxed{2}[\text{p-wrd}]) = \text{mkMtr}(\boxed{1}, \boxed{2})$$

$$(22) \text{extMtr}(\text{treasured}, \text{possession}) = \left[\begin{array}{c} \text{full-mtr} \\ \text{DOM} \quad \langle \text{treasured}, \boxed{1} \text{possession} \rangle \\ \text{DTE} \quad \boxed{1} \end{array} \right]$$

(ii) **The head contains post-head material** Perhaps the most awkward kind of mismatch between syntactic and prosodic structure arises when the complement or postmodifier of a syntactic head is ‘promoted’ to the level of sister of the constituent in which the head occurs; this creates a disjuncture between the lexical head and whatever follows. Fig 6 gives a typical example of this phenomenon, where the noun *possession* is followed by a prepositional complement, while Fig 7 represents the prosodic constituency. Let’s consider how *treasured* should combine with *possession of the samurai*. The Head-Complement Rule will have built a prosodic structure of the form *[possession [of the samurai]]* for the latter phrase. To obtain the correct results, we need to be able to detect that this is a metrical tree M whose leftmost element is a lexical head (by contrast, for example, with the structure *[treasured possession]*). In just this case, the extender can not only extend M but also create a new subtree by left-associating with the lexical head.² The required definition is shown in (23) and illustrated in example (24).

$$(23) \text{extMtr}(\boxed{1}[\text{full}], \left[\begin{array}{c} \text{DOM} \quad \boxed{2} \text{p-wrd} \oplus \boxed{3} \\ \text{DTE} \quad \boxed{4} \end{array} \right]) =$$

²The special prosodic status of lexical heads is incorporated in Selkirk’s (1981) notion of ϕ -phrase, and subsequent developments thereof, such as (Selkirk, 1986; Nespor and Vogel, 1986).

$$\left[\begin{array}{c} \text{full-mtr} \\ \text{DOM} \quad \text{extMtr}(\boxed{1}, \boxed{2}) \oplus \boxed{3} \\ \text{DTE} \quad \boxed{4} \end{array} \right]$$

provided that $\boxed{2}$ is the lexical head.

$$(24) \text{extMtr}(\text{this}, \left[\begin{array}{c} \text{full-mtr} \\ \text{DOM} \quad \langle \text{possession}, \boxed{1} [\text{of the samurai}] \rangle \\ \text{DTE} \quad \boxed{1} \end{array} \right]) = \left[\begin{array}{c} \text{full-mtr} \\ \text{DOM} \quad \left\langle \left[\begin{array}{c} \text{DOM} \quad \langle \text{this}, \boxed{2} \text{possession} \rangle \\ \text{DTE} \quad \boxed{2} \end{array} \right], \boxed{1} [\text{of the samurai}] \right\rangle \\ \text{DTE} \quad \boxed{1} \end{array} \right]$$

Turning back briefly to the Head-Specifier Rule shown in Fig 5, we can now see that if ϕ_0 is a metrical tree M , then the value of $\text{extMtr}(\phi_1, \phi_0)$ depends on the syntactic information associated with the leftmost element P of that tree. That is, if P is the phonology of the lexical head of the phrase, then it can be prosodically disjoined from the following material, otherwise the metrical tree M is extended in the standard way.

There are various ways that this sensitivity to syntactic role might be accommodated. One option would be to inspect the DTRS (daughters) attribute of a sign. However, I will briefly sketch the treatment implemented in the ALE grammar, which does not build a representation of daughters. Instead, I have introduced an attribute LEX inside the value of HEAD which is constrained in the case of lexical items to be token-identical to the PHON value. For example, the type for *possession* is approximately as follows:

$$(25) \left[\begin{array}{c} \text{word} \\ \text{PHON} \quad \boxed{1} \text{possession} \\ \text{SYNSEM} \quad \left[\begin{array}{c} \text{SYN} \mid \text{HEAD} \quad \left[\begin{array}{c} \text{noun} \\ \text{LEX} \quad \boxed{1} \end{array} \right] \\ \text{ARG-ST} \quad \langle \text{PP} \rangle \end{array} \right] \end{array} \right]$$

Since LEX is a head feature, it percolates up to any phrase projected from that head, and allows the PHON value of the lexical head to be accessed at that projection; i.e., headed phrases will also bear a specification $[\text{LEX } \textit{phon}]$, which can be interpreted as saying “my lexical head’s phonology value is *phon*”. In addition, we let the function extMtr in Fig 5 take as an extra argument the HEAD value of the mother, and then test whether the leftmost Prosodic Word in the metrical tree being extended is the same as the LEX value of the mother’s HEAD value.

(iii) **Extending the head with a Leaner** Finally, there is an additional clause to accommodate the case where the extending element is a Leaner. This triggers a kind of left association, in that the result of

combining *a* with [*treasured possession*] is a structure of the form $[[a \text{ treasured}] \text{ possession}]$.

$$(26) \text{ extMtr}(\boxed{1} \text{ lnr}, \begin{bmatrix} \text{DOM} & \langle \boxed{2} \rangle \oplus \langle \boxed{3} \rangle \\ \text{DTE} & \boxed{4} \end{bmatrix}) = \begin{bmatrix} \text{full-mtr} \\ \text{DOM} \text{ extMtr}(\boxed{1}, \boxed{2}) \oplus \langle \boxed{3} \rangle \\ \text{DTE} & \boxed{4} \end{bmatrix}$$

This will also allow an unaccented subject pronoun to left-associate with the lexical head of a VP, as in $[[\text{he provoked}] [\text{the objections of everyone}]]$ (Gee and Grosjean, 1983).

4 Concluding Remarks

I believe that the preceding analysis demonstrates that despite the well-known mismatches between syntactic and prosodic structure, it is possible to induce the required prosodic structures in tandem with syntax. Moreover, the analysis retains rather conventional notions of syntactic constituency, eschewing the non-standard syntactic constituents advocated by Prevost and Steedman (1993), Steedman (1990; 1991).

Although I have only mentioned two syntactic rules in HPSG, the radically underspecified nature of these rules, coupled with rich lexical entries, means that the approach I have sketched has more generality than might appear at first. With the addition of a rule for prenominal adjectives, prosodically interpreted like the Head-Specifier Rule, we can derive a range of analyses as summarised in (27). Here, I use square brackets to demarcate trees of type *full-mtr* and parentheses for trees of type *lnr-mtr*.

- (27) a. [this possession](of the samurai)
 b. [this treasured possession](of the samurai)
 c. (a treasured) possession
 d. (a treasured) possession [(of these) people]
 e. Kim gave (the book) (to the boy)
 f. Kim (gave it) (to the boy)
 g. Kim is happy [about Lee]
 h. Kim is happy [(that Lee) is fond (of the bird)]
 i. Kim wanted (to rely) (on the report) [(that Lee) is fond (of the bird)]

It would be straightforward to augment the grammar to accommodate post-modifiers of various kinds, which would behave prosodically like post-head complements. By contrast, auxiliaries do not conform to the association between headed structures and prosodic structures that we have seen so far. That is, if auxiliaries are a subtype of complement-taking verbs, as assumed within HPSG, then they depart from the usual pattern in behaving prosodically like specifiers rather than heads.

There are numerous directions in which the current work can be extended. In terms of empirical coverage, a more detailed account of weak function words seems highly desirable. The approach can also be tested within the context of speech synthesis, and preliminary work is underway on extending the Festival system (Black and Taylor, 1997) to accept input text marked up with metrical trees of the kind presented here. In the longer term, the intention is to integrate prosodic realisation within the framework of an HPSG-based concept-to-speech system.

Acknowledgements

I am grateful to Philip Miller, Mike Reape, Ivan Sag and Paul Taylor for their helpful comments on various incarnations of the work reported here.

References

- J. Bachenko and E. Fitzpatrick. 1990. A computational grammar of discourse-neutral prosodic phrasing in English. *Computational Linguistics*, 16(3):155–170.
- Mary E. Beckman and Janet B. Pierrehumbert. 1986. Intonational structure in English and Japanese. *Phonology Yearbook*, 3:255–310.
- Mary E. Beckman. 1986. *Stress and Non-Stress Accent*. Foris, Dordrecht, Holland.
- Steven Bird. 1995. *Computational Phonology: A Constraint-Based Approach*. Studies in Natural Language Processing. Cambridge University Press.
- Alan W. Black and Paul Taylor. 1997. The festival speech synthesis system. Technical Report TR-83, Human Communication Research Centre, University of Edinburgh, Edinburgh, UK, January.
- Bob Carpenter and Gerald Penn, 1999. *ALE: The Attribute Logic Engine. User's Guide*. Bell Laboratories, Lucent Technologies, Murray Hill, NJ, version 3.2 beta edition.
- Noam Chomsky and Morris Halle. 1968. *The Sound Pattern of English*. Harper and Row, New York.
- Steffan Corley, Martin Corley, Frank Keller, Matthew W. Crocker, and Shari Trewin. 1999. Finding syntactic structure in unparsed corpora: The gsearch corpus query system. *Computers and the Humanities*.
- James Paul Gee and François Grosjean. 1983. Performance structures: a psycholinguistic and linguistic appraisal. *Cognitive Psychology*, 15:411–458.

- Janet Hitzeman, Alan W. Black, Paul Taylor, Chris Mellish, and Jon Oberlander. 1998. On the use of automatically generated discourse-level information in a concept-to-speech synthesis system. In *ICSLP'98*, pages 2763–2768.
- D. Robert Ladd. 1996. *Intonational Phonology*. Cambridge University Press, Cambridge.
- Mark Liberman and Alan Prince. 1977. On stress and linguistic rhythm. *Linguistic Inquiry*, 8:249–336.
- Michael Mastroianni and Bob Carpenter. 1994. Constraint-based morpho-phonology. In *Proceedings of the First ACL SIGPhon Workshop*, Los Cruces, New Mexico. Association for Computational Linguistics.
- Marina Nespor and Irene Vogel. 1986. *Prosodic Phonology*. Number 28 in *Studies in Generative Grammar*. Foris Publications, Dordrecht.
- Janet B. Pierrehumbert and Mary E. Beckman. 1988. *Japanese Tone Structure*. Number 15 in *Linguistic Inquiry Monographs*. The MIT Press, Cambridge, MA.
- Carl Pollard and Ivan A. Sag. 1994. *Head-Driven Phrase Structure Grammar*. CSLI and University of Chicago Press, Stanford, Ca. and Chicago, Ill.
- Scott Prevost and Mark Steedman. 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, The Netherlands, April 21–23. OTS (The Research Institute for Language and Speech).
- Elisabeth Selkirk. 1981. On prosodic structure and its relation to syntactic structure. In T. Fretheim, editor, *Nordic Prosody II: Papers from a Symposium*. Tapir, Trondheim.
- Elisabeth O. Selkirk. 1984. *Phonology and Syntax: The Relation between Sound and Structure*. Current Studies in Linguistics. MIT Press, Cambridge, Mass.
- Elisabeth O. Selkirk. 1986. On derived domains in sentence phonology. *Phonology Yearbook*, 3:371–405.
- Stefanie Shattuck-Hufnagel and Alice E. Turk. 1996. A prosody tutorial for investigators of auditory sentence processing. *Journal of Psycholinguistic Research*, 25(2):193–247.
- Mark Steedman. 1990. Intonation and structure in spoken language understanding. In *Proceedings of the 28th Annual Meeting of the Association for Computational Linguistics*, pages 9–16, Pittsburgh, Pa., June. University of Pittsburgh.
- Mark Steedman. 1991. Structure and intonation. *Language*, 67(2):260–296, June.
- Arnold M. Zwicky. 1982. Stranded *to* and phonological phrasing in English. *Linguistics*, 20(1/2):3–57.
- Arnold M. Zwicky. 1986. The unaccented pronoun constraint in English. In Arnold M. Zwicky, editor, *Interfaces*, volume 32 of *Ohio State University Working Papers in Linguistics*, pages 100–114. Ohio State University Department of Linguistics, July.