# Prosodic resolution of syntactic ambiguity

Speech Technology term paper,
Fall semester 2004,
Lilja Øvrelid, Göteborg University,
lilja.ovrelid@svenska.gu.se

## 1 Introduction

This paper deals with the use of prosodic information in disambiguating syntactic structure, i.e. in arriving at the correct interpretation of a syntactically ambiguous sentence. We will start out by examining closer the linguistic phenomenon of prosody, its acoustic and functional properties. It is pretty clear that prosody plays an important role in human speech understanding, however, the exact nature of this role has been heavily investigated within the (psycho-)linguistic literature, and we will accordingly review some of these studies and attempt to arrive at an approximation regarding the main findings there.

Whether knowledge about prosody also can be utilized in automatic speech understanding is a relevant question in the present context. We will, as in Nöth et al. (2000), differentiate between automatic speech *recognition* (ASR) and speech *understanding* (ASU), where the former refers to the process of converting an acoustic signal into a string of words and the latter the application of further linguistic processing, such as syntactic parsing, to this string of words. For the purpose of this paper, the main focus will be on automatic speech understanding. This represents a difficult task, which is complicated further by syntactic ambiguity, as many other systems of natural language processing. We will therefore review some of the work done on incorporating prosodic information into syntactic parsing and by doing so attempting to provide a better understanding of the problems faced and the practical viability of such an incorporation within actual systems of speech understanding.

## 2 Prosody

Prosody describes aspects of speech which relate to units larger than phonemes, which makes it a *suprasegmental* property. Examples of these suprasegmental properties are pitch, loudness and duration. In the following section we will examine more closely some of these properties relevant to prosodic phenomena and their acoustic near-correlates.

The relative pitch of an utterance relates to its "melody". Its phonetic correlate is the "frequency of vibration of the vocal folds during the voicing of segments" (Laver, 1994, p. 450). Acoustically, this frequency of vibration is measured in cycles per second or Hertz, and determines the *fundamental frequency (F0)*. The range of F0-variation differs between the sexes, as well as between different speakers. The variation of pitch-span within an utterance is said to have a *pitch-contour*, describing the shape and direction of change in pitch, and is usually provided with a mnemonic name such as fall-rise, rise-fall-rise etc. *Intonation* is represented by changing patterns of F0 across prosodic phrases (i.e. across spans of speech greater than a phoneme or syllable).

Loudness relates to the physical concept of intensity, which correlates acoustically with the amplitude of oscillations of air molecules in soundwaves. These oscillations are measured in decibels (dB), which describes a relative scale centered around a reference sound.

Duration is related to the temporal, rhythmical organization of language material, and is therefore also related to breaks and pauses in speech.

# 3 Functions of prosody

The main function of prosody in speech is, in some sense, to structure the language material. Bruce (1998) cites three main areas of prosodic influence in speech: *prominence*, *grouping* and signalling aspects related to the *discourse*. Languages differ with regards to which functions dominate and the extent of prosodic usage, from tonal languages, where prosody is both semantically and syntactically crucial, to a language like English, where prosody figures mostly at the level of discourse.

## 3.1 Prominence

The notion of prominence is a complex one, involving on the physical side, higher pitch, greater loudness, longer duration and greater articulatory excursion (Laver, 1994). In speech, prominence or lack thereof can signal accentuation or deaccentuation of a segment or whole utterance. To take Swedish as an example, Bruce (1998) differentiates between three levels of increasing prominence: stress, accent and focus (focal accent). Stress influences the rhythmical aspect of an utterance and has a complex phonetic correlate involving at least intensity and duration, but probably also the spectral properties of the transition between vowel and consonant (Bruce, 1998). Accent and focus are mainly signaled by tonal means, which incorporate the rhythmical dimension (stress), as well as pitch, to signal increased prominence. Swedish is a tonal language and employ two word accents contrastively. Focus or focal accent signals the highest degree of prominence, usually by a longer duration of the focused element (Bruce, 1998).

Prosodic prominence has been claimed to be closely linked to the notion of the "information structure" of an utterance, reflecting the flow of information and the ordering of new vs. given information. A common generalization is that given information, i.e. information that has already been mentioned or is assumed by the speaker to be known, tends to precede new information. If, for some reason, this unmarked structuring of information is departed from, this is usually signaled prosodically, for instance by a focal accent. Another, related generalization, is that new information tends to be accented, whereas given information tends to be deaccented (Hirschberg, 2002). However, information structure alone is not not the sole instigator of accentuation. Terken and Hirschberg (1994) investigated the relationship between accent and the new vs.

given distinction in an experiment where the additional factors of syntactic function and sentence position were looked at. They found that when a given item has a certain syntactic function in the context of an utterance and this changes (for instance a subject in the context is an object in the present utterance), it is just as likely to be accented as an element which is new to the discourse.

### 3.2 Grouping

Grouping illustrates the intersecting natures of prosody and syntax and is employed both for prosodic phrasing i.e. indicating the boundaries of phrases, but also for coherence, i.e. indicating which parts of the utterance belong together. The relationship between syntax and prosody, is however, a non-isomorphic relationship. It is not necessarily the case that prosodic means are employed in order to group phrases. Bruce (1998) claims that prosody is more likely to be employed the "higher" in the syntactic tree the phrase in question is located. So, a sentence, for instance, will most likely be delimited by prosodic means, whereas a modifying prepositional phrase will not necessarily be so (although it certainly may).

Pauses are central to grouping at a higher syntactic level, e.g. between clauses, and temporal properties of speech in general are central to signaling phrasal boundaries, wherein shorter duration usually signals coherence and longer duration signals boundaries. The prosodic function of prominence, as reviewed above, is also involved in prosodic grouping. In particular, a rising F0 often signals the beginning of something new, hence the end of what preceded.

### 3.3 Discourse

Intonation is the key property of the discourse-related functions of prosody, but loudness also plays an integral part (Bruce, 1998). Some examples of these discourse-related functions of prosody include the signalling of speech acts, as for instance in the sentence below, which with a certain intonation functions as a yes/no question rather than a declarative:

(1) You have not seen the prime minister.

Signalling wanted feedback from a listener is another discourse-related function, where intonation and loudness contribute.

With regards to the main topic of this paper, prosody in syntactic disambiguation, we might conclude that the two proposed functions of prominence and grouping are the most relevant. By way of prosody speakers may for instance signal phrase boundaries, hence giving clues as to the syntactic interpretation of an utterance.

## 4 Prosody in human speech understanding and disambiguation

It is clear from the above that there exists a relationship between prosody and syntax, albeit a non-isomorphic one. Even so, the extent of this interplay - the use of prosody in speech production as well as its influence on human speech understanding - is an issue which is under debate in the literature.

Most of the studies that make claims regarding the role of prosody in human speech understanding and, in particular, in syntactic disambiguation, base these claims on experiments designed to elicit certain types of prosody. The level of “naturalness” of speech found in these experiments differ quite a bit, as do the results. Some studies make use of read speech - ambiguous sentences in disambiguating contexts, such as Hirschberg and Avesani (1997) and Weber et al. (2004), whereas several other studies make use of some sort of task-oriented dialog or role-play, with differing levels of predetermination of the speech (Snedeker and Trueswell, 2003; Bard and Aylett, 1999; Schafer et al., 2000). Most of these cite the necessity for studying spontaneous speech, rather than laboratory speech, however the actual “spontaneity” of the speech used in experiment situation is not always as realistic as one might wish.

In the following we will take a closer look at some of the claims made in the literature regarding the intricate relationship between prosody and syntactic disambiguation in human speech understanding.

## 4.1 Attachment ambiguities

The syntactic ambiguities receiving the most attention in the reviewed articles are so-called attachment ambiguities, e.g PP-attachment, relative clause attachment and attachment of other types of adverbials etc.

Hirschberg and Avesani (1997) compare Italian and English, and examine both syntactic and semantic disambiguation in recorded, read speech, where the ambiguous sentences are embedded in a completely disambiguating context. Among the syntactic ambiguities examined are PP- and adverbial attachment, as in (2a) and (2b) below:

(2) a. He managed to find the woman with the binoculars
    b. He had spoken to her quite clearly

They report no significant trends in prosodic patterning for the ambiguous PP-attachment sentences, but do report on increased consistency for the adverbial attachment and relative clause cases. Also, the sentences containing semantic ambiguities exhibit consistent prosodic patterning across most of the subjects, for instance in focus variation for the adverb *even* in sentences like (3) below, where the focus will be either on the verb *telegraphed* or on the object *the paper*:

(3) He even telegraphed the paper

In contrast, an experiment reported in Snedeker and Trueswell (2003), reports clear use of prosody for syntactic disambiguation. This study deals only with PP-attachment and involves a game-like task with a truly ambiguous setting (a set of items including a toy frog wearing a flower, one without a flower and a large life-size flower) and the following ambiguous sentences:

(4) a. Tap the frog with the flower (modifier interpretation)
    b. Tap the frog with the flower (instrument interpretation)

A following experiment with an unambiguous setting, i.e. a setting strongly favoring one of the readings, was also performed to control the results. The subjects of the experiments were given scripted versions of the sentences to use, but had to memorize and speak them at the appropriate time in the game, as chosen by themselves. The results of the first, main experiment, where the setting was ambiguous, were measured by the consequent action performed by the other participant in the game (i.e. either tapping the frog by using the flower or tapping the frog wearing the flower using his/her finger), indicating whether the sentence had been correctly disambiguated or not. Also, the sentences were analyzed and transcribed, [1] giving a fair idea of the relevant prosodic cues given to disambiguate. The listeners chose the correct interpretation in about 70% of the ambiguous cases in the main experiment, which is above chance, but lower than for the unambiguous "filler" sentences in the same experiment. The primary prosodic cue in the ambiguous sentence was duration and pausing. Snedeker and Trueswell (2003) conclude that there is a clear "ambiguity avoidance", wherein disambiguating prosody is only employed in settings where the sentence is not already disambiguated by the context or in other ways. This might explain the negative findings of Hirschberg and Avesani (1997) as well, as their ambiguous sentences where completely disambiguated by the linguistic context.

## 4.2 Local argument ambiguities

Schafer et al. (2000) investigate the late vs. early closure ambiguity caused by an optionally transitive verb as in (5a) and (5b) below (Schafer et al., 2000, p. 171):

(5) a. When that moves the square will..

b. When that moves the square it..

They also set up a game-like task for the subjects of the experiment, who work together two and two, with a predetermined set of sentences to choose from. The sentences are recorded and transcribed, and used in a comprehension test for new subjects in a second experiment, where only the ambiguous parts of the sentences in (5) were presented in a forced continuation experiment. It was from this experiment that the results were drawn, and these showed that the subjects were able to choose the correct continuation in about 74% of the cases, when provided only with the ambiguous part of the sentence, hence no linguistic context or other aids for disambiguation apart from prosody. Schafer et al. (2000) propose that negative results for the use of prosody in PP-attachment ambiguities (as in Hirschberg and Avesani (1997)) may be caused by the fact that these are adjuncts and not arguments as in the present study, where the choice is between a subordinate clause object reading, as in (5b), and a main clause subject reading, as in (5a) above. This is an interesting point worth investigating further. It might be the case that argument disambiguation in some sense is ranked as being more important than adjunct disambiguation, as arguments, after all, contribute more to the crucial content of the sentence.

Finally, a study which deals only with the disambiguation of arguments is reported in Weber et al. (2004). This article features an anticipatory eye-gazing experiment and deals with temporal ambiguities (that is, local ambiguities that are resolved later in the sentence), like Schafer et al. (2000) above. The ambiguous sentences are German main clauses where the first argument bears syncretic case marking which is ambiguous between nominative and accusative case, hence

[1] All of the studies reviewed in this section transcribed their data using the ToBI transcription standard or some augmented version of this, if for another language than English.

a subject and an object reading. Since German in addition may have both a SVO and an OVS word order, temporal ambiguity occurs when an ambiguously marked NP initiates a sentence (Weber et al., 2004, p. 7):

(6) a.

| | | | | |
|---|---|---|---|---|
| Die Katze | jagt | womöglicht | den | Vogel |
| the cat-NOM-amb. | chases | possibly | the | bird-ACC |

'The cat is possibly chasing the bird'

b.

| | | | | |
|---|---|---|---|---|
| Die Katze | jagt | womöglicht | der | Hund |
| the cat-ACC-amb. | chases | possibly | the | dog-NOM |

'The cat is possibly chased by the dog'

The ensuing experiment made use of recorded speech wherein a native speaker of German was made aware of the ambiguities and told to disambiguate them prosodically (but not how to do that, of course!). The recorded speech was then played to a subject who was watching a picture of several animals, but with no actions between these. The subject's eye movements were registered and thought to be indicative of the sentence interpretation of the subject at that point in time, relating the expectations (hence 'anticipatory') of the listener when hearing only the ambiguous parts of the sentence. The results showed that there were clear prosodic differences in the SVO and OVS sentences in the recorded speech. In SVO sentences, like (6a), the nuclear stress was on the main verb, whereas in the OVS sentences, like (6b), it was on the initial object. The anticipatory eye movements following the verb were in a significant number of the cases in accordance with the prosodic cues, indicating that their interpretation of the local ambiguity was resolved through prosodic means.

There are however, some problems with this study, when compared with the above. First of all, the experiment made use of recorded speech where the speaker had been made aware of and told to disambiguate the ambiguous structures. Also, the fact that the animals in the picture were known entities, probably influenced the intonation and created an even sharper prosodic difference between the SVO and OVS sentences, as the authors themselves propose. However, the fact that the listener seemed to be able to make use of the prosodic cues in disambiguation is an interesting result of the study and indicates that the use of prosody is not arbitrary, but serves a function in the interface with syntax.

As we have seen, the results from experiments relating prosody and syntactic ambiguity are varying, to say the least, and they certainly confirm our earlier statement that there is a complex, non-isomorphic relationship between the two. We have also seen that the data and experimental set-ups employed vary quite a bit, making the comparison between different studies difficult. It does, however, seem clear that the human subjects were able to distinguish and make use of disambiguating prosody in the cases where that was employed by the speaker, whether in completely ambiguous settings or not. The types of ambiguities dealt with also vary, however, the majority deal with attachment ambiguities. These are also the most relevant as we now move into the realm of automatic speech understanding, where local ambiguities like the ones presented in section 4.2 are obviously not a problem.

# 5 Prosody and syntactic disambiguation in ASU

Making use of prosodic information in automatic speech understanding seems to be a natural consequence of the fact that there is a relationship between prosody and syntax. As we have seen in the preceding section, however, this relationship is complex in nature. In the following we will take a closer look at how prosodic information has been utilized in automatic speech understanding, and examine some of the literature - the proposals made and the obtained results - in order to better evaluate the usability of prosody in speech understanding, and, in particular, in syntactic disambiguation.

Utilizing prosodic information in syntactic disambiguation will minimally involve a two-step process: (i) extraction of relevant prosodic features, and (ii) putting these features to actual use in a syntactic parser of some kind. In the following we will take a look at some proposals as to how this two-step process may be constructed and what the results are in terms of syntactic disambiguation.

So, what kind of ambiguities are typically dealt with in the literature? The syntactic ambiguities that are prevalent are mostly similar types of ambiguities as in the psycholinguistic experiments described above. These typically include prepositional attachment and preposition vs. particle ambiguities, which are notoriously difficult to resolve in NLP. These ambiguities are thought to be (at least partially) resolved by prosody, and in particular, the prosodic functions of grouping and prominence.

As in the linguistic studies dealing with prosodic ambiguity resolution, much of the work in automatic speech understanding involves using a corpus of read speech, containing certain syntactic ambiguities. Bear and Price (1990) make use of 35 such sentences, containing mostly PP-particle ambiguities. They introduce the notion of *break indices* between words in a sentence, numerically indicating the level of attachment between two words, where the lower the number is, the "tighter" the attachment. An example taken from Bear and Price (1990) is provided below, where the break indices are in place. Note the higher number between the verb and the preposition for the PP-reading:

(7) a. The 0 men 1 won **3** over 0 their 0 enemies (PP-reading)
b. The 0 men 2 won **0** over 1 their 0 enemies (particle reading)

Bear and Price (1990) base their automatic assignment of these word break indices on normalized aspects of the prosodic property of duration. In order to incorporate the prosodic information into their rule-based grammar, additional categories are added to the rules in order to allow for break indices between words, and also, constraints on the possible values of these indices in certain syntactic positions were made. The results obtained when applied to the 35 ambiguous sentences were mixed - whereas the number of readings from the parser decreased with 25%[2], the efficiency deteriorated and parse times increased by 37%.

Veilleux et al. (1993) also base their disambiguation on break indices assigned automatically, however, the assignments do not represent abolutes, but

[2]Actually, the disambiguation of PP vs. particle as described in Bear and Price (1990) was only applicable in the cases where a break index was a large number, i.e. in PP-readings in this case, because "a small prosodic gap does not provide a reliable way to tell which two constituents combine" (Bear and Price, 1990, p. 20).

rather probabilities of being followed by a certain break index, and this information contributes in a process of statistical parse scoring along with other lexical and syntactic factors. The correct parse is then equated with the most probable parse, given a range of factors. The results reported by Veilleux et al. (1993) are somewhat better than the ones from Bear and Price (1990). Most importantly, they result in an improvement on the disambiguation of both readings of a sentence like (7) above. For the PP-reading they achieve 80% correct disambiguation, but for the particle reading the performance is only slightly above chance, at 55%. However, Veilleux et al. (1993) clearly show that the problem resides in the automatic break index assignment, because when the parser is tested with manually assigned breaks, the performance improves drastically from 80 to 100 percent for the PP-reading and from 55 to 95 percent for the particle reading.

Whereas Bear and Price (1990) employ prosody in syntactic disambiguation by adding additional information to the rules of the parser's grammar, Veilleux et al. (1993) rather employ a post-parsing method, where the prosodic scores for each parse is computed after the syntactic parsing of the sentence. In contrast, Potisuk et al. (1996) use prosodic information as an additional source of knowledge, available in parallel with knowledge on syntactic constraints. They make use of a constraint dependency parser and this additional prosodic knowledge in order to disambiguate Thai sequences ambiguous between a compound reading and a verb-noun reading. This ambiguity is thought to be resolved prosodically by *stress*, which, as we remember from above, has a complex phonetic correlate. The task of stress classification is therefore an important part of this study, prior to the actual prosodic-syntactic disambiguation may take place. They employ a Bayesian classifier which makes use of five acoustic features represented in a feature vector. The classifier is trained on read speech containing ambiguous sentences in a disambiguating environment, and used to classify word hypothesis graphs[3] from the recognizer. Unfortunately, Potisuk et al. (1996) do not provide any solid results from their study, which makes it difficult to compare with the others.

All of the above studies, make use of read speech devised to disambiguate a fairly limited range of sentences and constructions within small, task-specific systems. Nöth et al. (2000) and Batliner et al. (2001), however, describe the use of prosody in a large-scale speech-to-speech translation system, namely that of the Verbmobil project. They employ prosody in many of their modules, syntactic and semantic analysis, dialog processing, transfer and speech synthesis. In terms of feature extraction, Nöth et al. (2000) make use of a highly redundant set of features (as many as 276 different features!), collected in a feature vector. Rather than deciding on the relevant feature set before extraction, they make use of a redundant feature set and let the statistical classifier determine which features are relevant. They make use of mostly acoustic-prosodic features (aspects of duration, F0-variation, intensity/energy length of pauses etc.), but also some lexical features, dealing with word accents. Hence each word hypothesis graph is enriched with this information and classified with regards to boundary strength and accent. The enriched graphs representing sentence hypotheses are then parsed using a unification-based parser. As mentioned above, the input

[3]A word hypothesis graph is a directed acyclic graph, where "each edge corresponds to a word hypothesis which has attached to it its acoustic probability, its first and last time frame, and a time alignment of the underlying phoneme sequence" (Nöth et al., 2000, p. 526)

to the system is in fact spontaneous speech and not read speech as in the previous studies. The results obtained for the syntactic parser are therefore quite extraordinary: "the statistics show that on average, the number of readings decreases by 96% when prosodic information is used, and the parse time drops by 92%" (Nöth et al., 2000, p. 528). However, due to the nature of their input, the grammar of their parser really takes seriously the nature of spontaneous speech, with all its possibilities for ellipses, extra-positions etc. Therefore the original, pre-prosodic grammar contains vast amounts of ambiguity for which prosody is a crucial disambiguator.

# 6 Discussion and conclusion

In the above sections we have examined how prosody contributes to the process of speech understanding, and in particular, to syntactic parsing. We have seen that prosody seems to play a role in human speech understanding, and that we are able to make continuous use of prosodic information in the process of syntactic disambiguation. We have also seen that prosody may be utilized as a source of information in systems for automatic speech understanding, most successfully, as we have seen, in the recent Verbmobil project.

There are several difficulties in the work with prosody, both in linguistic and computational linguistic respects. First of all, and as has already been stated, the relationship between prosody and syntax is a complex one. Although the use of prosody serves many functions in human communication, it is almost never a compulsory means to communicate something, i.e. there are many ways to say very similar things, wherein prosody only serves as one possibility (Hirschberg, 2002; Bruce, 1998). Also, and conversely, prosody serves many functions and "the same prosodic feature can be used to communicate many different meanings" (Hirschberg, 2002). In the psycholinguistic experiments performed above, isolation of one function of prosody was attempted through a controlled environment with limited linguistic influence. In the inclusion of prosody into automatic speech understanding, and specifically in the work of Nöth et al. (2000), the attitude was rather one of letting the statistics speak for themselves, i.e. bombarding the algorithm with a myriad of prosodic features and letting it make the appropriate generalizations based on the actual speech data. The fact that this procedure was the one with most success, speaks for the mere complexity of the phenomenon. According to Batliner et al. (2001), the main direction for the development of the use of prosody in automatic speech understanding may be summarized under the slogan *Prosody goes multi!* We have already seen that the approach of Nöth et al. (2000) was *multi-feature*, but Batliner et al. (2001) also emphasize *multi-level* and *multi-knowledge*, i.e. prosody alone is not enough but should be combined with other levels and sources of linguistic information. This interplay relates to the proposed "ambiguity avoidance" in human speech understanding (Snedeker and Trueswell, 2003), where prosody is claimed to be used mainly in cases where other linguistic sources have not already disambiguated the utterance.

Finally, there is also the *multi-functional* property of prosody, as mentioned above, which should be taken seriously. The results from linguistic experiments have been somewhat inconclusive, as to the production of disambiguating prosody and a plausible consequence of that on the computational side would be to make use of probabilistic, soft, rather than hard constraints in parsing: “we do not make hard decisions based on prosodic events in order to prune the search space. We rather guide the search in the huge search space by using probabilities about prosodic events” (Nöth et al., 2000).

## References


E.G. Bard and M.P. Aylett. The dissociation of deaccenting, givenness, and syntactic role in spontaneous speech. In J. Ohala, Y. Hasegawa, M. Ohala, D. Granville, and A. C. Bailey, editors, *Proceedings of the XIVth International Congress of Phonetic Science*, pages 1753–1756, 1999.

Anton Batliner, Elmar Nöth, Jan Buckow, Richard Huber, Volker Warne, and Heinrich Niemann. Whence and whither prosody in automatic speech understanding: A case study. In *Proceedings of the ISCA tutorial and research workshop on speech recognition and understanding*, pages 23–28, Red Bank, NJ, 2001.

John Bear and Patti Price. Prosody, syntax and parsing. In *Proceedings of the 28th annual ACL*, 1990.

Gösta Bruce. *Allmän och svensk prosodi*, volume 16 of *Praktisk Lingvistik*. Institutionen för lingvistik, Lunds Universitet, 1998.

Julia Hirschberg. Communication and prosody: functional aspects of prosody. *Speech Communication*, 36:31–43, 2002.

Julia Hirschberg and Cinzi Avesani. The role of prosody in disambiguating potentially ambiguous utterances in English and Italian. In *Proceedings of the ECSA Tutorial and research workshop on intonation*, Athens, 1997.

John Laver. *Principles of phonetics*. Cambridge University Press, 1994.

Elmar Nöth, Anton Batliner, Andreas Kießling, Ralf Kompe, and Heinrich Niemann. VERBMOBIL: The use of prosody in the linguistic components of a speech understanding system. *IEEE Transactions on speech and audio processing*, 8(5), 2000.

Siripong Potisuk, Mary P. Harper, and Jack Gandour. Using stress to disambiguate spoken Thai sentences containing syntactic ambiguity. In *Proc. ICSLP '96*, volume 2, pages 805–808, Philadelphia, PA, 1996. URL `citeseer.ist.psu.edu/507626.html`.

Amy J. Schafer, Shari R. Speer, Paul Warren, and S. David White. Intonational disambiguation in sentence production and comprehension. *Journal of psycholinguistic research*, 29(2):169–182, 2000.

Jesse Snedeker and John Trueswell. Using prosody to avoid ambiguity: Effects of speaker awareness and referential context. *Journal of memory and language*, 48:103–130, 2003.

Jacques Terken and Julia Hirschberg. Deaccentuation of words representing ‘given’ information: effects of persistence of grammatical function and surface position. *Language and Speech*, 37(2):125–145, 1994.

N. Veilleux, M. Ostendorf, and C. Wightman. Parse scoring with prosodic information. In *Proc. Int. Conf Acoustics, Speech, Signal Processing*, volume 2, pages 51–54, Minneapolis, MN, 1993.

Andrea Weber, Martine Grice, and Matthew W. Crocker. The role of prosody in the interpretation of structural ambiguities: a study of anticipatory eye movements. Submitted, 2004.