Computational Modelling of Contextual Coreference : Implications for Swedish Text-to-Speech

Merle Horne*, Marcus Filipsson*, Mats Ljungqvist‡ and Anders Lindström‡

* University of Lund, ‡Telia Promotor Infovox

Abstract

The issue of describing and identifying coreference relations between lexical (content) words is discussed. In text-to-speech applications, it is important to be able to recognize these relations since contextually coreferent (given) words are associated with tone accent patterns that differ from those on new information. An algorithm for referent tracking in a restricted domain is also described which allows one to preprocess a text and automatically tag words as either contextually 'New' or 'Given'. The algorithm presupposes computational modelling of lexical semantic identity of sense relations as well as information on inflexional/derivational morphology and compounding. This information is available in a lemmatized lexicon of Swedish. Referent identity is defined on head-word representations derived from the text input on the basis of the inflexional expansion rules contained in the lexicon. Information on the New/Given status of words can subsequently be used in the F_0-generating component of the text-to-speech system to trigger the assignment of focal vs non-focal word accents.

I INTRODUCTION

Researchers on prosody have been aware for some time of the consequences the coreferential ('contextually given') status of words has on sentence intonation (Danes* 1967, Gunter 1966). In phonetic investigations, the 'new/given' parameter has been used by e.g. Bruce (1977) in order to trigger focal accents on different words in test sentences. It has also been applied in a model for assigning levels of sentence stress in English (Horne 1988).

One of the principal goals in current research in text-to-speech has involved improving the prosody component by developing linguistic preprocessors which extract such contextual information. Recent work on the development of interfaces for English text-to-speech systems has begun to show the results of efforts to take into account contextual factors that condition intonational patterning. Hirschberg (1990) for example has attempted to implement fragments of Grosz & Sidner's (1986) model of discourse by modelling aspects of their 'attentional state', in particular the notion of a 'stack' of focus spaces. These have been represented by a stack of morpheme roots ('lemmas') which is updated at certain fixed intervals, e.g. at paragraph boundaries. Words figuring in this stack are thus not accented by the text-to-speech rules. Youd & House (1991) also keep track of mentioned roots in an 'accessibility table' in their work on modelling information in the restricted area of a telephone dialogue system for flight information. Our own research efforts have been directed towards further developing modelling of the different kinds of contextual information which have implications for the generation of intonation in Swedish.

II IDENTITY RELATIONS

One of our endeavours has involved computationally modelling the different kinds of contextual phenomena that can affect a particular word's status in a text with respect to the parameter 'new/given' (Horne et al. 1993). This involves being able to track, in addition to relations based on morphological identity such as in the work of Hirschberg and Youd & House mentioned above, more complex identity of sense relations that are used in situations of anaphora in order to refer back to an already-mentioned item, e.g. synonymy, hyponymy and part-whole (meronymy) relations (Allerton 1978, Lyons 1977, Miller & Fellbaum 1991). All these different types of 'cohesive devices' (Halliday & Hasan 1976, Morris & Hirst 1991) that are used in situations of anaphora in order to express textual givenness (coreference) are

important to model for purposes of intonation generation since they lead to the assignment of nonfocal accents on anaphoric lexical words. In summary, these anaphoric devices include the following (identical subscripts indicate coreference; words in uppercase letters are focussed):

1) **Reiteration using a morphologically identical form:**
E.g. *Marie skulle vilja ha en SAAB$_i$, men inte en RÖD Saab$_i$.*
'Marie would like to have a SAAB$_i$, but not a RED Saab$_i$.'

2) **Reiteration using a previously mentioned morpheme with a different inflexional/derivational ending:**
E.g. *Marie tycker att vi borde investera i en SAAB$_a$ men jag måste erkänna att jag faktiskt inte TYCKER OM Saabar$_a$,*
'Marie thinks that we should invest in a SAAB$_a$, but I must admit that I really don't LIKE Saabs$_a$'.

3) **Reiteration by means of a synonym:**
E.g. *Ingvar tror att en skattehöjning är MÖJLIG$_k$, men Carl säger att en sådan åtgärd inte är på något SÄTTgenomförbar$_k$.*
'Ingvar thinks that a tax-rise is POSSIBLE$_k$, but Carl says that such a measure is not at ALL feasable$_k$'.

4) **Reiteration by means of a superordinate term:**
a) **Hyponymy:**
E.g. *Min son vill ha en TAX$_m$, men jag är inte säker på att han är gammal nog att ta HAND om en hund$_m$.*
'My son wants a DACHSHUND$_m$, but I am not sure he is old enough to take CARE of a dog$_m$'.

b) **Meronymy (part/whole):**
E.g. *Min HÅRDDISK$_i$ kraschade igen imorse. Det måste vara den SÄMSTA dator$_i$ jag någonsin HAFT!*
'My HARD DISK$_i$ crashed again this morning. It's got to be the WORST computer$_i$ I've ever HAD!'

Note that in the case of hyponymy and meronymy relations, the identity relation is non-symmetrical, so that it is essential that the more general term follows the more specific term in order for the latter to be interpreted as an anaphor to the former and to be assigned a non-focal accent (see Allerton 1978). Placing the more specific term (e.g. *tax* 'dachshund') after the more general term (e.g. *hund* 'dog') does not trigger a non-focal accent, as the inappropriate accentual pattern in (5) illustrates:

5) **Min SON vill ha en HUND, men jag är inte SÄKER på att han är GAMMAL nog att ta HAND om en tax'.*
'*My SON wants a DOG, but I'm not sure he's OLD enough to take CARE of a dachshund'.

One reason why text-to-speech systems have not implemented detailed semantic rule systems for keeping track of the above mentioned identity of sense relationships is that in order to be able to handle unrestricted texts where one has no specific knowledge of the context, the amount of structuring that would be required and the corresponding computational time involved in searching through all the possible semantic relationships would be horrendous. To take just one rather simple illustrative example, the word *skiva* 'disk' is polysemous, and consequently has many semantic associations to many other different words depending on the context in which they are used. In the context of physiology, for example, it refers to a part of the human body; in the context of computers, it refers to a component used for storing information; in the context of phonographs, it also refers to an object for storing information, in

this case, sound; however, in this context, it is related to other words such as *platta* 'record', *singel* 'single', *LP-skiva* 'album' that a computer *disk* is not associated with at all. Thus, because of this semantic ambiguity, in order to decide whether the word *skiva* 'disk' is coreferent with some other lexical item that was previously mentioned in a text, it would be necessary to search through all the possible semantic relationships that all the meanings of *skiva* can have with other lexical items. Modelling this kind of encyclopedic knowledge is not currently feasible. Consequently, in order to develop a tractable computational model for structuring lexical information, it seems more reasonable to attempt to represent knowledge in a more limited domain where the ambiguities associated with polysemy can in most cases be avoided.

III SEMANTIC RELATIONS IN A SPECIFIC DOMAIN: STOCK MARKET

One domain that lends itself to study in this respect is the stock market which is a very well-defined world involving buying and selling transactions in stocks and bonds. The reports that appear daily in the newspapers for example deal with changes that the system undergoes during a given day. One important lexical structure that is important for determining coreference relations in stock market reports is constituted by the relations between individual companies and their superordinate structures, such as the partial hierarchy shown in (6):

(6)

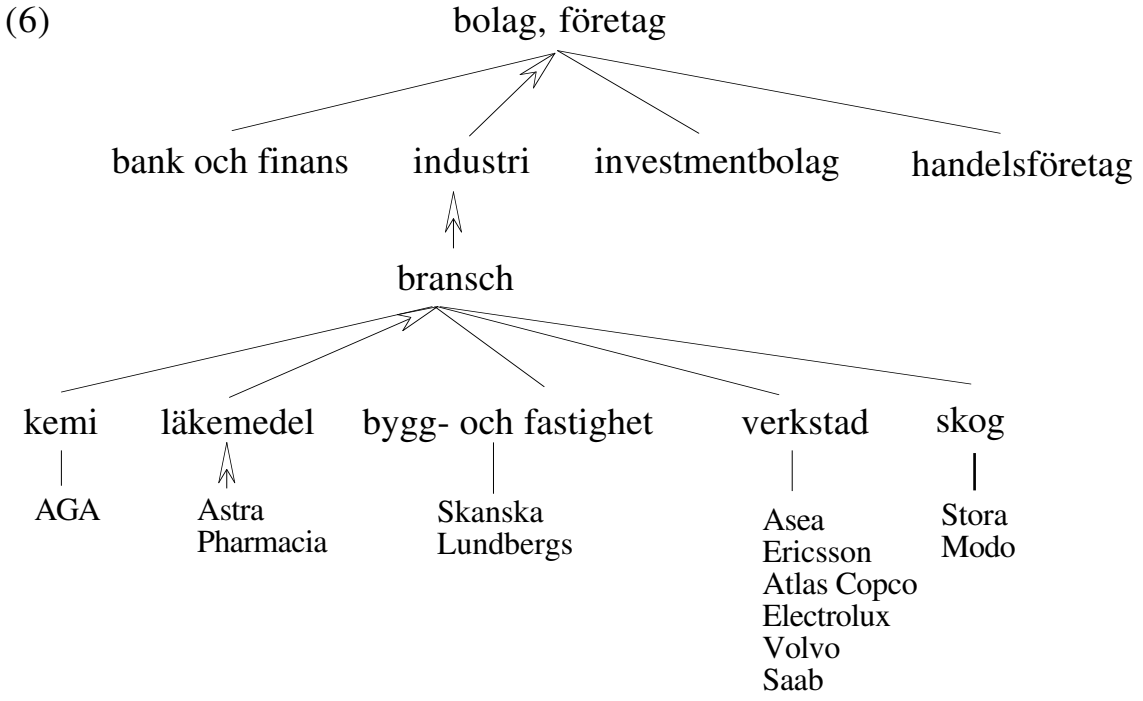

This hierarchy represents a partial classification of individual share names into more general classes. For example, it indicates that *Astra* is a *läkemedel* 'drug' company, and that *läkemedel*, *bygg-och fastighet* 'construction and real estate', and *skog* 'forest' are all *branscher* 'branches' of *industri* 'industrial companies' which in their turn are grouped together as *bolag* 'companies' together with other more specific groups of companies such as *bank och försäkring* 'bank and insurance', *investmentbolag* 'investment companies', etc. Thus, one would expect that for example, after having mentioned *Astra* shares, one could refer to them anaphorically by calling them shares in *läkemedel* 'drugs'.

Another hierarchical semantic relation that is present in the stock market world is that realized in expressions for *aktier* 'stocks' and *obligationer*.'bonds' These can be referred to anaphorically by the superordinate term *papper* 'paper' One can also see a rather rich inventory of terms used to cite the value of a given stock on the market; there would seem to be a relation

of synonymy between the terms *index* 'index', and *kurs* 'rate', and a relationship of hyponymy between them and the more general term *nivå* 'level' These can be represented as in (7):

(7)

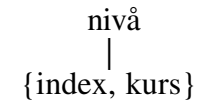

One can also expect that certain concepts in the stock market world are 'situationally/pragmatically given' (i.e. taken for granted) for all speaker/hearers (see Chafe 1974, Firbas 1979). A candidate for this status is the word denoting the particular currency used for a given country's stock exchange reports (in Sweden, the *krona* 'crown'). One can thus mark such words in the lexicon as pragmatically given and consequently, they will not trigger focal accents. The same goes for situational information such as the day and year of the particular stock market report since this information is 'given' in the newspaper.

IV STOCK MARKET DATA

In order to gain more detailed knowledge of the lexical structure of stock market texts, we made a preliminary lexical and prosodic investigation of a newspaper text taken from *Sydsvenska Dagbladet*. The text was read and recorded by a male native speaker of Standard Swedish. A portion of the text is reproduced below in (8). Coreference relationships have been marked by means of identical bold subscripts. The symbol '+' is used in the Swedish text to show the boundary between component morphemes in a compound.

(8) Stockholms**a** fond+börs**b** general+index**c** slutade**d** på torsdagen**e** på 858,8, en uppgång**f** med marginella 0,02 procent**g** jämfört med onsdagens**h** slut**d**+index**c**. Kurs**c**+utvecklingen**i** över dagen**e** betecknades som oregelbunden.

Kurs**c**+stegringar i AGA**j** och Astra**k** fick bransch**j,k**+index**c** för kemi**j**- och läkemedel**k** att bli bästa bransch**j,k** med en uppgång**f** på 1,6 procent**g**.

Förlorare var övrig industri**j,k** och handels+företag, som föll med 2,6 respektive 1,9 procent**g**.

Stockholm's**a** stock exchange**b** general index**c** closed**d** on Thursday**e** at 858.8, a marginal increase**f** of 0.02 percent**g** compared with Wednesday's**h** closing**d** index**c**. The rate**c** development**i** during the day**e** was described as irregular.

Rate**c** increases**f** in AGA**j** and Astra**k** made the trade index**c** for chemicals**j** and drugs**k** to be the best line of business**j,k** with an increase**f** of 1.6 percent**g**.

The losers were the remaining industrial and trading companies**j,k**, which fell by 2.6 and 1.9 percent**g**, respectively.

V PROSODIC REFLEXES OF NEW/GIVEN PARAMETER IN SWEDISH

Although Swedish is a Germanic language, its prosodic system is considerably different from e.g. English and German since, unlike these latter, it is a 'pitch accent' language. It has two lexical word accents termed Accent 1, or Acute accent, which can be represented as HL*, and Accent 2, or Grave accent, which can be represented as H*L (Bruce 1977). In Standard Swedish, there is, furthermore, a very salient difference between focal and non-focal accent patterns. When focused, the word accents are followed by a H tone, viz. HL*H (Accent 1) and H*LH (Accent 2). In addition to the absence of the final H tone, a non-focal Accent 1 or 2 is usually characterized by a relatively narrower register than a focal one. Even compound words are characterized by the absence of a final H tone in non-focal contexts, i.e. the compound word accent pattern H*L...L*H, where H* is associated with the first stressed syllable and L* to the last stressed syllable becomes H*L...L* in non-focal contexts. In Figures 1a and 1b

(adapted from Bruce 1977) are schematized Swedish focal and non-focal accent patterns in simple and compound words:

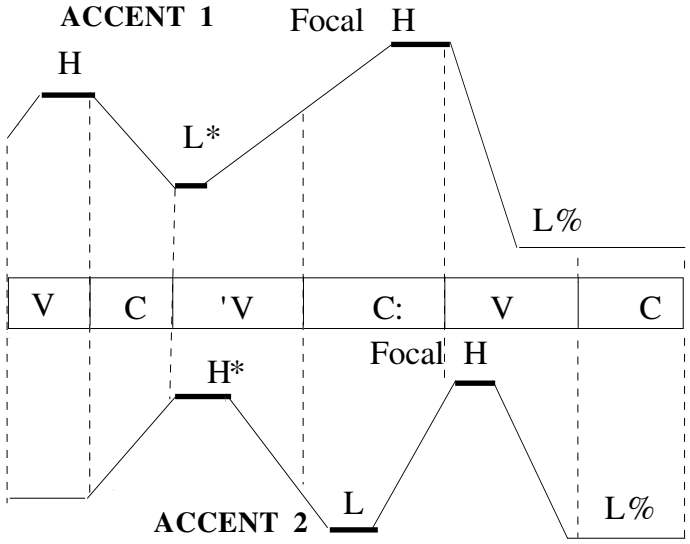

Figure 1a. Schematic representation of Swedish word accents. The association of the starred tone with the stressed vowel ('V) is critical. The other associations are only approximate (see Bruce 1977).

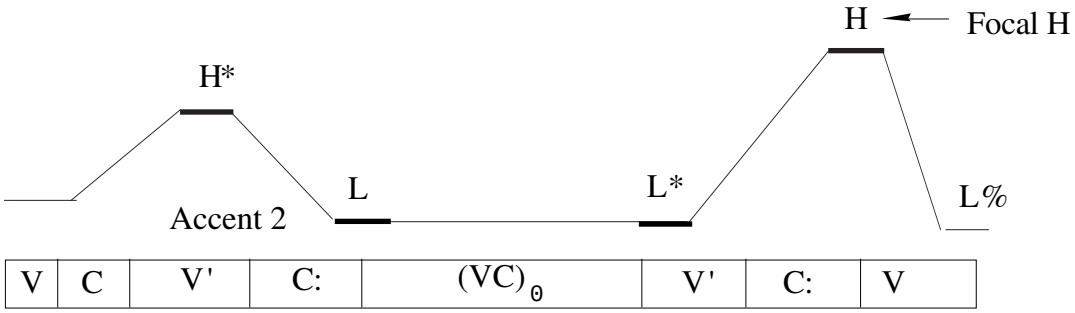

Figure 1b. Accent patterning in Swedish compound words. Association of the starred tone to the stressed vowels ('V) is critical. The other associations are approximate and can vary (see Bruce 1977)

VI OBSERVED TONE-ACCENT PATTERNS IN STOCK-MARKET DATA

As one can see from examining the newspaper text above in (8), there are a considerable number of anaphoric relations present. The second and third paragraphs contain instances of the superordinate structure in (6). In this text, one observes that the specific companies *AGA* and *Astra* are referred to later in the text by *kemi och läkemedel* 'chemical and drug (companies)' These latter terms are then subsequently referred to by the expression *bransch* 'branch', and in the final sentence, the expression *industri* 'industry' is used to refer back to the individual branch industries. Thus, following the argumentation above, we would expect that *AGA* and *Astra* would be assigned focal accents, since they constitute new information; however, we would not expect that the superordinate terms that are used to refer back to them anaphorically, i.e. *kemi och läkemedel*, *bransch*, *industri* would be assigned focal accents. This is, in fact, what one finds when one examines Fo patterns on the words under discussion. In Figure 2 are presented the Swedish intonational patterns associated with *AGA och Astra*, *kemi och läkemedel* and *bransch.*

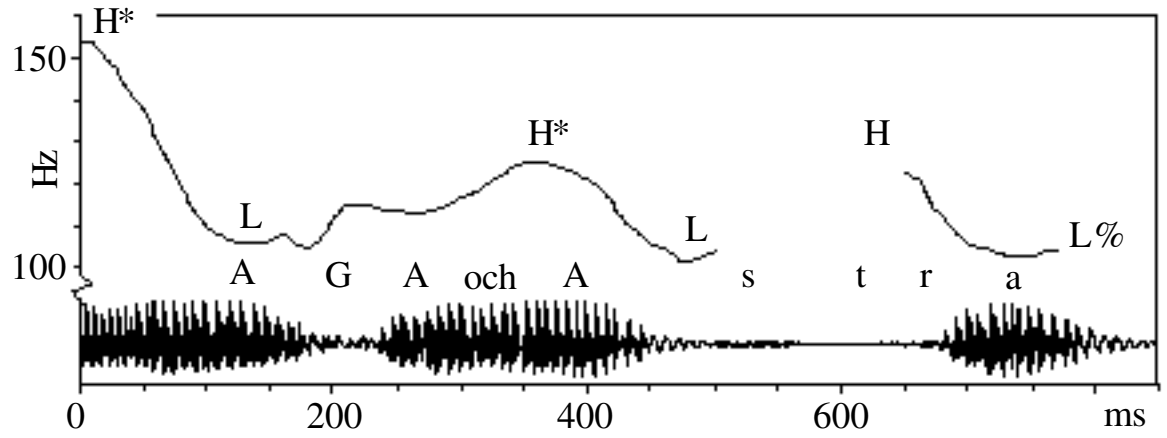

Figure 2a. Fo pattern associated with the phrase *AGA och Astra* ([+new]). Both words are assigned Accent 2 (H*L). *Astra* is also assigned a 'focal' H in its final syllable.

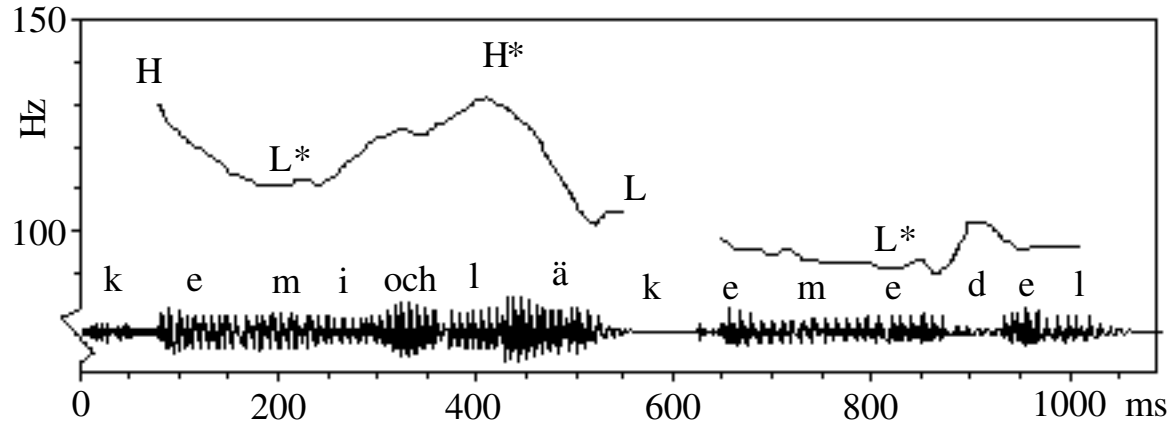

Figure 2b. Fo pattern on *kemi och läkemedel* 'chemical and drug (companies)' ([-new], i.e. given information). *Kemi* exhibits a non-focal HL* accent pattern as does the compound word *läkemedel*, i.e. it has an Accent 2 (H*L) pattern on the first stressed syllable, and a L* on the last stressed syllable.

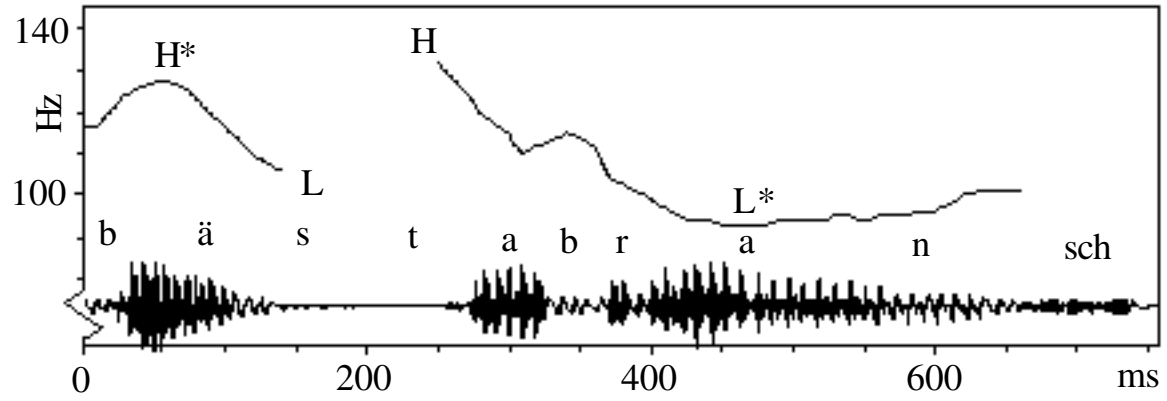

Figure 2c. Fo pattern on *bästa bransch* 'best branch' where *bästa* constitutes [+new] information and *bransch* [-new] information. *Bästa*, an Acc. 2 (H*L) word, is thus associated with a focal H in its final syllable. *Bransch*, an Acc.1 word is not, however, associated with a focal H.

VI COMPUTATIONAL MODELLING OF NEW VS GIVEN INFORMATION

In most commercial text-to-speech systems, the prosodic component does not have access to higher level syntactic and semantic information and is therefore able to generate only a very restricted number of intonation patterns. In the case of Swedish (Carlson & Granström 1973, Bruce & Granström 1989), the default rules assign a focal accent to the last content word in an intonational phrase. This focal accent pattern leads to the interpretation of phrase-final words as 'new' information. It has thus not been possible to *automatically* assign phrase-final content words non-focal accents which are associated with 'given' information. In Standard Swedish, however, there is, as we have shown above, a very salient difference between focal and non-focal word accents. In order to enable automatic assignment of the proper accent-type, it is necessary to process the input text with respect to the information content associated with the lexical words. We have currently been involved in developing such a linguistic preprocessor which models and tracks coreferential relationships between content words. In this endeavour, we have exploited the information contained in a computerized lexicon of Swedish (Hedelin et al. 1987). The lexicon, which is lemmatized, contains approximately 116,000 headwords, each one listed with its part-of-speech specification, inflection code, and phonetic transcription. It also includes information on the morphological status of derivations and can handle compound-words, either by explicit listing or by algorithmic generation. This is a necessity in languages like Swedish and German where compounds are written without indication of internal word boundaries. In Fig. 3 are presented a number of entries illustrating the general structure of the lexicon.

omstrukturera	{v1}	/ˆåmm~struktu^{2}re:ra/	<Ssng om#1A+strukturera>
omstrukturering	{s2}	/ˆåmm~struktu^{2}re:riñ/	<Avldng omstrukturera+-ing>
onsdag	{s11}	/'onsta/	<>
oregelbunden	{a1}	/ˆo:~^{1}re:gël^{2}bundën/	<Avldng o-+regelbunden>
ovanlig	{a1}	/ˆo:~^{2}va:nlig/	<Avldng o-+vanlig>
ovanligt	{adv}	/ˆo:~^{2}va:nlit/	<Avldng ovanlig+-t>
på#A	{adv}	/'på:/	<>
på#B	{prep}	/'på:/	<>

Figure 3. Illustrative entries from Hedelin et al.'s lexicon showing some of the information available. Accompanying each head-word is {word category specification} (*s* = substantiv 'noun'), /pronunciation/ (including word-accent type) as well as information concerning the morphological make-up of compounds (*Ssng*) and the structure of derived forms (*Avldng*).

A general feature of a lemmatized lexicon is the inherent relationships between the head-words and their inflected forms. As an example, the word *tända* 'to light' can be related to its paradigmatic forms: *tända* (inf.), *tänder* (pres.), *tände* (pret.), *tänt* (past part.), *tänd* (supine), *tändas* (inf. passive), etc. This is important in the present application since the identity relations are defined over the head-words. The lexical structure has furthermore been amended with domain-specific knowledge describing semantic hierarchical relations (hyponymy, part/whole) and synonymy relationships.

The input text is analyzed in the following manner: First, the lexicon is used to analyse the words and decompose them into morphs (see Fig. 4). As mentioned earlier, the lexicon handles inflexions, derivations, and compound words. The treatment of compounds is an important feature, since the referent-tracking processes must apply not only to the compound as a whole, but to the component morphs. In Fig. 5 is presented an example of the decomposition of words into morphs.

The next step involves checking each word for its possible coreference with any previously mentioned word within an adjustable window. The window used in the examples presented in this paper has somewhat arbitrarily been chosen to be 60 words, but other domains could also be considered, e.g. the paragraph (Hirschberg 1990). The referent-tracking algorithm consists of four parts, as shown in Figure 4 (a-d):

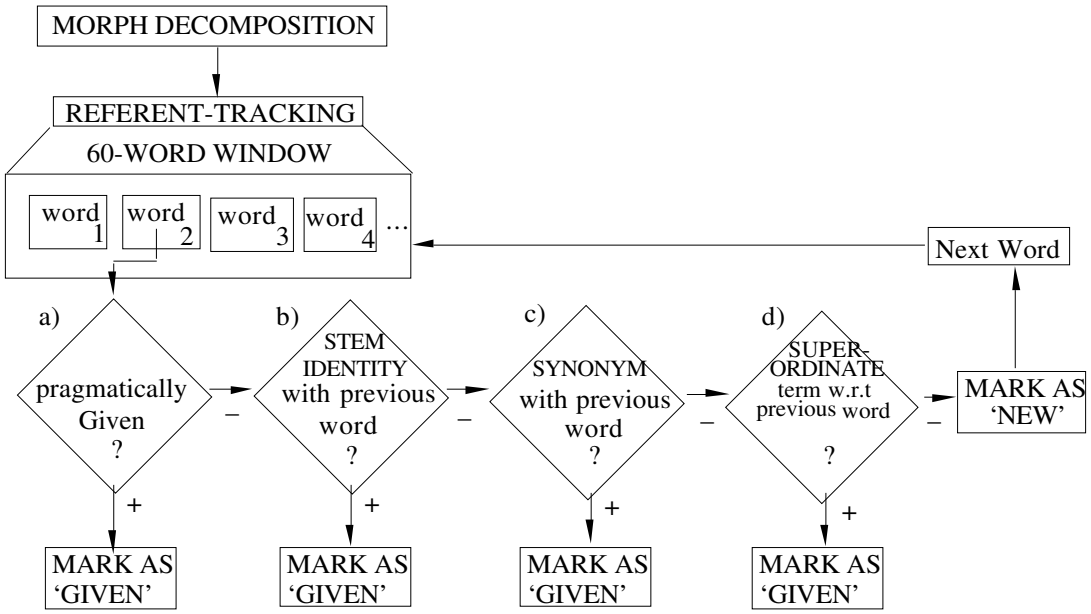

Figure 4. The different stages of the analysis in the lexical processor: MORPH-DECOMPOSITION first analyses each word into its content morphs. In the second stage, the REFERENT-TRACKING algorithm searches through a 60-word window for coreferents.

The first one tracks and marks words that are situationally/pragmatically 'given', such as *börs* 'stock-market', and *krona* 'crown' in the domain we have studied. The second stage involves identifying cases of coreference due to reiteration of root morphs (or stems), as obtained from the initial analysis using the lexicon. The third part uses 'domain-specific' synonymy relations to identify cases of coreference. Examples from the stock-market domain are: *kurs–nivå* 'rate' and *aktie–papper* 'share'. The fourth and final stage attempts to track words that are involved in superordinate identity relations (hyponymy, meronymy relationships). In order to do this, hierarchical relations such as those in Figure 6 have been modelled using 'is an example of' or 'is a part of' pointers to establish the relation between pairs of lemmas thus building up a forest of hierarchical, multi-branch trees. Cases of multiple inheritance, i.e. where a lemma has more

:Stockholms	'Stockholm's'	:torsdagen	'Thursday+def.art'	:marginella	'marginal+def'
Stockholm	'Stockholm'	torsdag	'Thursday'	marginal	'marginal'
:fondbörs	'stock-exchange'	:på'at'		:0_02	'0.02'
fond	'stock'	på	'at'	0_02*	'0.02'
börs	'exchange'	:858_8	'858.8'	:procent	'percent'
:generalindex	'general-index	858_8*	'858.8'	procent	'percent'
general	'general'	:en'an'		:jämfört	'compared'
index	'index'	en	'an'	jämföra	'compare (inf.)'
:slutade	'closed'	:uppgång	'increase'=going up	:med	'with'
sluta	'close (inf.)'	upp	'up'	med	'with'
:på	'on'	gå	'go (inf.)'	:onsdagens	'Wednesday's'
på	'on'	:med	'with'	onsdag	'Wednesday'
		med	'with'	:slutindex	'closing-index'
				sluta	'close (inf.)'
				index	'index'

Figure 5. Example of decomposition of words into morphs using a lemmatized lexicon. Analysed words are preceded by a colon (:). Under each of these words is/are the component root-morph(s). Notice that for verb-forms, the infinitive form is taken to be the root. An asterisk following a form indicates that there is no corresponding lexical entry.

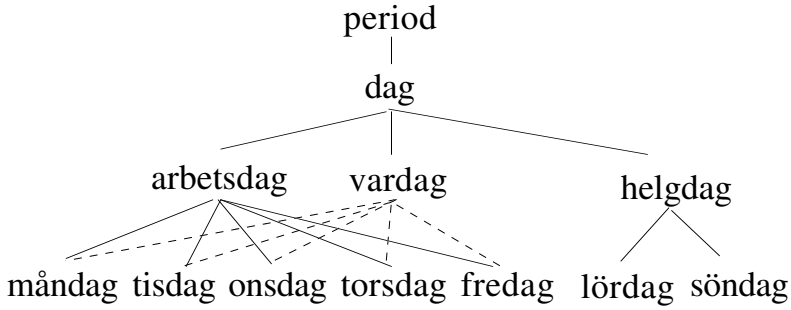

Figure 6. Example of a superordinate hierarchy modelled in the lexical processor. A word in the hierarchy is marked as 'Given' if it is preceded in the text by a more specific word, i.e. one that occurs lower down in the tree.

than one parent (e.g. both *vardag* 'week-day' and *arbetsdag* 'work-day' are parents of *måndag* 'Monday') do not pose a problem, since the algorithm searches only among 'daughters' of the currently analyzed word, i.e. the tracking in semantic hierarchies is unidirectional, from the more general anaphor to the more specific antecedent. Output from the text processing scheme as described above is shown in Figure 7. Each word is marked as either 'new' (N) or 'given' (G).

1**N**		Stockholms	16**G**i	11	med	31**G**i	3	branschindex
2**G**g		fondbörs	17**N**		onsdagens	32**N**		för
3**N**		generalindex	18**G**i	3	slutindex	33**G**h	27	kemi
4**N**		slutade	19**N**		Kursutvecklingen	34**G**w		och
5**N**		på	20**N**		över	35**G**h	29	läkemedel
6**N**		torsdagen	21**G**h	6	dagen	36**G**w		att
7**G**i	5	på	22**N**		betecknades	37**N**		bli
8**N**		858_8	23**N**		som	38**N**		bästa
9**N**		en	24**N**		oregelbunden	39**G**i	31	bransch
10**N**		uppgång	25**G**s	10	Kursstegringar	40**G**i	11	med
11**N**		med	26**N**		i	41**G**i	9	en
12**N**		marginella	27**N**		AGA	42**G**i	10	uppgång
13**N**		0_02	28**G**w		och	43**G**i	5	på
14**N**		procent	29**N**		Astra	44**N**		1_6
15**N**		jämfört	30**N**		fick	45**G**i	14	procent

Figure 7. An example of the output from the referent-tracking algorithm. Words are marked either as N for 'new' or G for 'given'. If a word is marked as G, the output shows in addition the kind of givenness that is present: Gg stands for pragmatically given, Gi, for 'given' due to morphological identity. The number directly preceding the word refers to the number of the preceding word to which the word is construed as coreferent. Gh stands for 'given' due to a hierarchical relationship, e.g. word 22, *dagen* is marked as Gh since it refers back to word 6, *torsdagen*. Gs refers to coreference due to synonymy, e.g. the component *stegring* 'rise' in the compound word 25 *kursstegringar* 'rate-increases' is a synonym to word 10 *uppgång*, thus triggering a non-focal accent on the compound. Gw stands for grammatical words that are not eligible to receive focal accents.

VII CONCLUSION

The linguistic processor described in the present paper has been implemented as a pre-processor to the Infovox text-to-speech synthesis system. It analyses the input text and annotates it with respect to 'new' and 'given' information, thereby allowing the system to generate synthetic prosody of considerably improved naturalness as compared to the default behaviour, as verified

by informal listening tests.The performance of the processing can, however, be improved by further analysis, such as identification of prosodic word and phrase boundaries (Horne 1994) as well as relative levels of pitch prominence. Information on the new/given status of lexical words is of course a prerequisite to generating more detailed prosodic structure since focussed words play a role in triggering prosodic phrase boundaries. Information on contextual coreference is also needed in order to predict certain cases of contrastive prominence, in particular, those involving parallel syntactic structures as in (9) (see Horne 1987) :

9) John's dachshund$_i$ nipped$_j$ Tim's son$_k$ and then MARY'S dog$_i$ bit$_j$ JOE'S little boy$_k$.

Thus referent tracking and modelling lexical semantic relations constitute analyses which are essential prerequisites to carrying out many other kinds of linguistic preprocessing in text-to-speech systems as well as in other natural language processing applications.

REFERENCES

Allerton, D. (1978), 'The notion of 'givenness' and its relations to presupposition and theme', *Lingua* 44: 133-168.

Bruce, G. (1977) *Swedish Word Accents in Sentence Perspective*, Gleerups, Lund.

Bruce, G. and B. Granström, (1989), 'Modelling Swedish intonation in a text-to-speech system', *STL-QPSR* 1: 31-36.

Carlson, R. and B. Granström (1973), 'Word accent, emphatic stress, and syntax in a text-to-speech system', *STL-QPSR*, 2-3: 31-36.

Chafe, W. (1974), 'Language and consciousness', *Language* 50: 111-133.

Danes*, F. (1960), Sentence intonation from a functional point of view', *Word* 16: 34-54.

Firbas, J. (1979), 'Post-intonation-centre prosodic shade in the modern English clause', in S. Greenbaum, G. Leech, and J. Svartvik (eds.), *Studies in English Linguistics for Randolph Quirk*, London, Longmans, 125-133.

Grosz, B. and C. Sidner (1986), 'Attention, intentions and the structure of discourse', *Computational Linguistics* 12: 175-204.

Gunter, R. (1966), 'On the placement of accent in dialogue: A feature of context grammar', *Journal of Linguistics* 2: 159-179.

Halliday, M. and R. Hasan (1976), *Cohesion in English*, Longmans, London.

Hedelin, P., A. Jonsson and P. Lindblad (1987), *Svenskt uttalslexikon*: 3 ed. Tech. Report, Chalmer's Univ. of Technology.

Hirschberg, J. (1990), 'Using discourse context to guide pitch accent decisions in synthetic speech'. *Proceedings, ESCA Workshop on Speech Synthesis*, Autrans, France, 181-184.

Horne, M. (1987), 'Contrastive prominence' and syntactic parallelism', in *Towards a Discourse-Based Model of English Sentence Intonation*, *Working Papers* (Dept. of Ling., Univ. of Lund) 32.

Horne, M. (1988), 'Towards a quantified focus-based model for synthesizing English sentence stress', *Lingua* 75: 25-54.

Horne, M. (1994), 'Generating prosodic structure for synthesis of Swedish intonation', *Working Papers (Dept. of Linguistics, Lund Univ.)* 43: 72-75.

Horne, M., M. Filipsson, M. Ljungqvist, and A. Lindström (1993), 'Referent tracking in restricted texts using a lemmatized lexicon: implications for generation of prosody'. *Proceedings Eurospeech '93* (Berlin) Vol. 3, 2011-2014.

Lyons, J. (1977), *Semantics*, Cambridge, Cambridge University Press.

Miller, G. and C. Fellbaum (1991), 'Semantic networks of English', *Cognition* 41: 197-229.

Morris, J. and G Hirst (1991), 'Lexical cohesion computed by thesaural relations as an indicator of the structure of text', *Computational Linguistics* 17: 21-42.

Youd, N. and J. House (1991), 'Generating intonation in a voice dialogue system', *Proceedings Eurospeech 91*, Genua, Italy, 1287-1290.