Chapter 4

Intonational Phrasing

4.1 Introduction

Our research objective is to design, build, and test a linguistically-based SLG system that generates speech with pitch accents and intonational phrases that are most acceptable to experts (Section 1.3). This chapter addresses intonational phrasing, i.e. the way the flow of speech is segmented into chunks by intonational means. These chunks have been referred to by terms such as *intonational domains*, *intonational phrases*, *tone units*, *tone groups*, and *association domains*. For current purposes, we will use the general term intonational phrase, which is defined as a stretch of speech covered by a perceptually coherent intonation contour. The boundaries of an intonational phrase correspond to a perceived discontinuity in the intonation contour, resulting from of a change in pitch level. In addition, they may be, and usually are, signalled by durational means such as pause or preboundary lengthening. In our view, however, a prosodic boundary marked by durational cues only does not qualify as an intonational boundary (Section 2.3.7).

The work described in this chapter does not directly concern SLG itself. However, since intonational phrasing is an essential task in any SLG system, an adequate approach to intonational phrasing is a prerequisite. Our goal here is to obtain an intonational phrasing algorithm for Dutch that can generate the preferred intonational boundaries for an utterance. This algorithm can then be incorporated in our SLG system, which will be the topic of the following chapters. An additional question in this respect is what linguistic information the algorithm requires as input, because this is the information that the NLG part of the SLG system will have to supply.

Although a considerable amount of linguistic theory about intonational phrasing is available in the literature, most work is based on a small set of crucial observations, and has neither been tested against larger sets of data (Gussenhoven 1992a; Hirst 1993; Ladd 1986; Nespor and Vogel 1986; Selkirk 1984), nor systematically compared with competing approaches (Altenberg

1990; Bachenko and Fitzpatrick 1990; Croft 1995). In the first part of this chapter, therefore, three intonational phrasing algorithms are systematically compared using a substantial set of reference data. The corpus-based method for evaluating these algorithms will be described in Section 4.2; the corpus itself in Section 4.3. In Section 4.4 to Section 4.6, we will successively describe each algorithm and analyse its errors. An intermediary discussion and conclusion will be provided in Section 4.7. Next, a new algorithm will be proposed and tested in Section 4.8. We will end this chapter with a summary.

4.2 Method

The evaluation method used here is essentially the same as that used in the previous chapter (Section 3.4.1) to evaluate phonological phrasing, and is a practical implementation of several methodological choices. As explained in Section 1.1.3, acceptability judgements regarding intonational phrasing may vary, because (1) they may be gradual rather than categorical, (2) alternative phrasings may be equally acceptable, and (3) intersubjective differences may occur. To cope with this variation, we have chosen to aim for the intonation that is most acceptable according to the judgements of intonational experts, ignoring differences between equally acceptable forms. In addition, we argued in Section 1.1.4 that the disadvantage of a perceptual evaluation is that other aspects of the wave form, such as segmental errors, inadequate phonetic realization, or speaker characteristics, may interfere with the judgement of the intonation. We therefore decided to rely on a symbolic assessment of the acceptability judgements of the experts, asking them to prescribe the preferred intonational form for a given texts and comparing this with the prediction by our algorithms. Accordingly, the evaluation method used here encompasses the following steps.

1. **Construction of the corpus**
   Assemble a corpus of sentences. For each sentence, add its syntactic/semantic structure and its preferred intonational structure, i.e. pitch accents[1] and intonational phrase boundaries.

2. **Implementation of intonational phrasing algorithms**
   Implement the intonational phrasing algorithms.

3. **Evaluation of intonational phrasing algorithms**

   (a) *Automatic comparison*
   Apply each algorithm to the sentences in the corpus. For each sentence, automatically compare the predicted boundaries to the preferred boundaries. Report all differences and generate a confusion matrix.

[1] Some of the intonational phrasing algorithms make use of the pitch accent distribution.

(b) *Error analysis*
On the basis of the results of the automatic comparison, trace the origin of systematic errors in the algorithms. Find the advantages and disadvantages of each algorithm, and determine which algorithm has the best overall performance.

The implemented algorithms were selected on the basis of the following criteria. First, given our goal of building a linguistically-based SLG system, purely probabilistic approaches to intonational phrasing (Hirschberg 1991; Ostendorf and Veilleux 1994; Sanders and Taylor 1995; Taylor and Black 1998) were not taken into account.[2]. Second, the algorithm had to be applicable to Dutch, as this is the language we are primarily interested in. The majority of the algorithms address English, and are either not intended for other languages, or would require significant modifications when used for other languages. This excludes for example Altenberg (1990), and Bachenko and Fitzpatrick (1990). Third, the algorithm had to be explicitly described and sufficiently formalised in order to be implemented, which rules out proposals such as those by Dirksen and Quené (1993), Hirst (1993), and Selkirk (1984, 1995). Using these criteria, two algorithms were chosen that represent quite different views on intonational phrasing. The first is by Nespor and Vogel (1986, ch. 7), which is probably the best known algorithm for intonational phrasing, and the second by Gussenhoven (1990, 1992a, 1992a). We also looked at a revised version of the first algorithm as proposed by Caspers (1994, ch. 2). In addition, we tested a new algorithm.

All implementations were written in GRAMTSY (*Grammatical Transformational System*), a programming system that supports the use of transformational rules and grammars to manipulate trees in the form of labelled bracket structures (Coppen 1990). As far as the first three algorithms are concerned, lack of space makes it impossible to discuss all of them in depth, so we refer to the original publications for details. However, some of the problems that arose during implementation are mentioned in the footnotes.

4.3 Corpus

4.3.1 Text

The text material was the same as that used for the evaluation of phonological phrasing (Section 3.4.2), and consisted of 10 descriptions of wild flowers taken from a botanical handbook (Forey 1992). Each description consisted of 4 to 8 sentences. The total number of sentences was 52, the total numbers of words 907, an average of 17.4 words per sentence. An example is given in (1); a number of complete flower descriptions are given in Appendix A and Appendix B. Since most sentences were relatively long and have a rather complex syntactic structure, they formed an appropriate challenge for intonational phrasing

[2] However, these will be addressed in Section 8.5.3

algorithms. Their length and complexity gave rise to a large number of non-trivial instances of intonational phrasing. In contrast, the discourse structure of the texts was fairly simple, and essentially consisted of the introduction of a particular flower and an elaboration on its properties, minimising the potential influence of discourse structure on intonational phrasing was minimal.

(1) De Reuzenbalsemien is uit de Himalaya naar
The Impatiens-Glandulifera has-been from the Himalayas to
Europa gebracht en koloniseert hier waterkanten in vele
Europe brought and colonises here watersides in many
gebieden ook plaatselijk in Nederland en Belgie.
areas also locally in the-Netherlands and Belgium

4.3.2 Intonational structure

The reference intonational phrases were produced by the same panel that was responsible for the accent distribution used in the evaluation of phonological phrasing, and consisted of three intonationally experienced, native speakers of Dutch. They worked on a preprocessed version of the texts, in which commas were removed, semicolons[3] were changed to full stops, and abbreviations and numbers were expanded to their full lexical forms. They were asked to prescribe the intonational structure to be used in a competent reading of the text. This included the *accent distribution* and the *boundary distribution*, i.e. the locations of pitch accents and intonational phrase boundaries. As expected, boundaries turned out to always coincide with word junctures. A distinction was made between *obligatory* and *optional* boundaries, indicated by a '#' and a '|' respectively. This distinction was introduced to accommodate equally acceptable alternative phrasings (Section 1.1.3). The panel discussed disagreements in their phrasings until a consensus was reached. In total, they assigned 109 obligatory boundaries[4] and 42 optional boundaries, against 704 word junctures where no boundary was specified. The prescribed phrasing for (1) is given in (2); see Appendix A and Appendix B for some other sentences.

(2) De Reuzenbalsemien | is uit de Himalaya naar Europa gebracht #
en koloniseert hier waterkanten in vele gebieden # ook plaatselijk in
Nederland en Belgie. #

4.3.3 Syntactic structure

The sentences were manually parsed by the author according to the syntactic analysis outlined in Section 2.6.[5] These syntactic structures provide all infor-

[3] There were eight of these, each connecting a pair of main clauses, i.e. clauses with the finite verb in second position.

[4] We ignore the 52 trivial obligatory boundaries at the end of a sentence.

[5] We experimented with using syntactic parsers, but they either failed to handle the level of complexity of these sentences, or did not provide the required level of detail.

mation that is required by the intonational phrasing algorithms. The syntactic structure corresponding to the sentence in (2) was shown in Figure 3.2. A number of extra features were added to the parse trees to satisfy the needs of particular phrasing algorithms. For instance, to distinguish arguments from modifiers, arguments were linked to their corresponding predicates by means of a numerical index (cf. arg_1 and $pred_1$ in Figure 3.2). Furthermore, subjects, root sentences[6] (cf. *root* in Figure 3.2), non-restrictive relative clauses and appositional elements (cf. *appos* in Figure 3.2) were indicated by special features. Following Nespor and Vogel (1986), we used a broad definition of *apposition*, one which includes parenthetical expressions, non-restrictive relative clauses, tag questions, vocatives, expletives, and certain moved elements, not all of which necessarily occurred in the corpus.

4.3.4 Prosodic structure

All three intonational phrasing algorithms require information about the prosodic structure. $\mathcal{N}\&\mathcal{V}$ and $\mathcal{CAS}$ both need a parse of the utterance in terms of phonological phrases (φ's). For this, we used the φ-phrasing algorithm for Dutch as proposed in Section 3.3.3.1. In addition, $\mathcal{GUS}$ assumes as input a *recursive* φ-structure. The φ-structure on top of the base-level φ's is constructed in such way that the dominance relations among the φ's mirrors the dominance relations among the syntactic constituents they are derived from. Figure 4.1 shows the recursive prosodic structure derived for sentence (1).[7]

4.4 Intonational phrasing according to Nespor & Vogel

4.4.1 $\mathcal{N}\&\mathcal{V}$ algorithm

The intonational phrasing algorithm by Nespor and Vogel (1986, ch. 7) ($\mathcal{N}\&\mathcal{V}$) is the mapping rule for π's[8], which constitutes a part of their version of prosodic phonology (Section 2.2). In general, such rules receive input from two sources. The first is the morpho-syntactic structure of an utterance. As morphological distinctions are irrelevant for deriving π's, $\mathcal{N}\&\mathcal{V}$ only requires the syntactic structure in which root sentences and appositional elements can be identified. Second, a partially-derived prosodic structure is required, because it is assumed that the mapping into prosodic constituents that are lower on the prosodic hierarchy has already been accomplished. Therefore, $\mathcal{N}\&\mathcal{V}$ requires a parse of

[6] A *root sentence* is an S node that is dominated by S nodes only. The notion of *root sentence* stems from a period in transformational generative syntax in which *root transformations*, such as *Verb Second* and *Topicalisation* in Dutch, only applied to certain cycles. These were called *root transformations* and *root sentences*, respectively.

[7] The π's in this structure are required by $\mathcal{GUS}$ in order to derive AD's. This point will be discussed in Section 4.6.1.

[8] Nespor & Vogel use the symbol 'I' to designate an intonational phrase. We use π for reasons explained in Section 2.3.7.

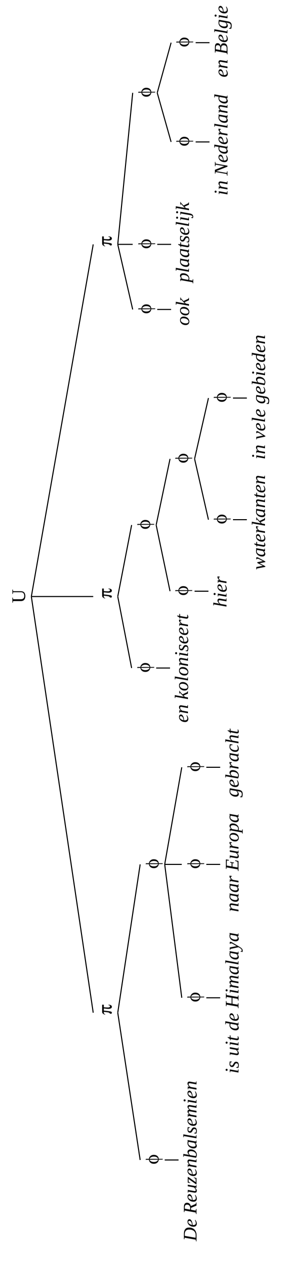

Figure 4.1: Prosodic tree for the sentence in (1); the prosodic categories are Utterance (U), Intonational Phrase (π), and Phonological Phrase (φ).

the utterance in terms of φ's as well. Its output consists of the locations of optional and obligatory π-boundaries.

The phrasing algorithm proceeds in two steps. In the first step, called *π-formation*, the initial π-boundaries are established.

(3) **π-formation**

Join into an n-ary branching π:

a. all the φ's in an appositional string;

b. any remaining sequence of adjacent φ's in a root sentence.

The next step, called *π-restructuring*, involves optionally dividing π's into smaller π's by adding boundaries. π's may be restructured due to factors such as the length of a π, speaking rate, style or contrastive prominence. Although the presence of such boundaries may be optional, their location is highly restricted by the syntax. According to the following definition, restructuring can only occur at certain syntactically defined positions.

(4) **π-restructuring** (optional)

a. Restructure after an NP if this does not interrupt another NP and does not separate an internal argument from its verb.[9]

b. Restructure before an S if this does not interrupt another NP.

The first part inserts optional boundaries after each argument of a verb, except the one that is adjacent to the verb. The second part will insert optional boundaries before an argument S, but not before a restrictive relative clause.

In addition, there is a special provision for restructuring lists.

(5) **List restructuring** (optional):

In a sequence of more than two constituents of the same type, i.e. $x_1, x_2, \ldots x_n$, restructure before each repetition of the node X, i.e. before $x_2, x_3, \ldots x_n$.

Applying $\mathcal{N}\&\mathcal{V}$ to the sentence in (1) gives the phrasing in (6), where ' # ' stands for an obligatory boundary and ' | ' for an optional boundary. Notice that there is no optional boundary after the NP *Europa*, as this would separate the verb *gebracht* from its internal argument *naar Europa.* Also, no boundary occurs after the NP *Nederland*, as this would interrupt the dominating NP.

(6) De Reuzenbalsemien | is uit de Himalaya | naar Europa gebracht # en koloniseert hier waterkanten in vele gebieden # ook plaatselijk in Nederland en Belgie. #

[9]As noted in Caspers (1994, p. 181), it is not clear whether a subject is an internal or obligatory argument. We follow Caspers who, based on example (18a) in Nespor and Vogel (1986, p. 197), concludes that subjects are not internal arguments. We will also assume that the end of an NP embedded in another NP does not constitute an optional boundary, unless an intervening S node occurs (Caspers 1994, p. 181).

Table 4.1: Confusion matrices for $\mathcal{N}\&\mathcal{V}$ algorithm. The top matrix contains the absolute number of confusions. The bottom matrix contains the percentage of confusions relative to the total number of boundaries in a column.

		Reference			
		Obl	*Opt*	*Non*	*Tot*
	Obl	39	0	0	39
N&V	*Opt*	42	34	50	126
	Non	28	8	654	690
	Tot	109	42	704	855

		Reference		
		Obl	*Opt*	*Non*
	Obl	35.78	0	0
N&V	*Opt*	38.53	80.95	7.10
	Non	25.69	19.05	92.90
	Tot	100	100	100

4.4.2 $\mathcal{N}\&\mathcal{V}$ results

4.4.2.1 General

Table 4.1 summarises the results for $\mathcal{N}\&\mathcal{V}$. The kappa value calculated for the matrix (without the optional boundaries) indicates that the agreement between the prediction and the reference is very low, and not significant at the 1%-level ($\kappa = 0.022$, $p = 0.045$). The most salient finding is that the algorithm predicts no incorrect obligatory boundaries. Apparently, all appositional elements and non-restrictive relative clauses are indeed delimited by an obligatory boundary in the reference. The other side of the coin is that, due to this conservative strategy for obligatory π's, 42 (38.53%) of the obligatory boundaries in the reference are predicted to be merely optional, whereas 28 of them (25.69%) are not predicted at all.

Looking at the predicted optional boundaries, we see that $\mathcal{N}\&\mathcal{V}$ generates a rather large number of incorrect optional boundaries. There are 50 optional boundaries at locations where no boundary should occur, and 42 optional boundaries where instead an obligatory one is required.

In the next subsections, we will look in more detail at the these errors in order to identify their causes. The reference prescribes both optional and obligatory boundaries, the two types that are predicted by the phrasing algorithms as well. Including the category 'absence of a boundary', this means

Table 4.2: Locations that were obligatory boundaries in the reference, but no boundaries according to $\mathcal{N}$&$\mathcal{V}$.

Syntactic configuration:	*Number:*
between AP's	8
before a PP	
before an attributive PP	6
before a non-attributive PP	3
between conjuncts	
between two NP conjuncts	6
between two VP conjuncts	2
before a restrictive relative clause	3
Total:	28

there are nine possible combinations of reference and prediction. To facilitate the discussion, boundary locations will be indicated as 'Reference/Prediction', where both can take the value of '#' (an obligatory boundary), '|' (an optional boundary) or '∅' (no boundary). Thus for example #/| denotes a location that is an obligatory boundary according to the reference, but an optional boundary according to the algorithm. Whenever a comparison between the two is irrelevant for the discussion, only the prediction will be mentioned.

If the reference prescribes an optional boundary, this is interpreted as: a boundary is acceptable here, but no boundary is equally acceptable. At such positions, therefore, whether the algorithm predicts a boundary, an optional boundary, or no boundary at all is irrelevant. In other words, where an optional boundary is allowed, the prediction can never be wrong. As a consequence, we will ignore incorrect predictions of optional boundaries of the types |/∅ and |/# , in the error analysis. Instead, we focus on important deviations of the types #/∅ , #/| , ∅/# , and ∅/| .

4.4.2.2 Obligatory in the reference, but none according to $\mathcal{N}$&$\mathcal{V}$ (#/∅)

Table 4.2 lists all syntactic configurations in which a boundary was obligatory according to the reference, but $\mathcal{N}$&$\mathcal{V}$ failed to predict one. Eight of these boundaries occurred between AP's. Although in general we find no boundary between two attributive AP's, a boundary becomes obligatory if the AP's are sufficiently long and/or complex. For example, the reference requires an obligatory boundary between the two AP's in (7), *groenachtig witte* and *netvormig geaderde*. However, $\mathcal{N}$&$\mathcal{V}$ is incapable of predicting a boundary at this position, since it is not the end of a NP, neither the start of an S, nor located between two listed items.

(7) . . . # en vijf groenachtig witte #/∅ netvormig geaderde
and five greenish white clathrate veined
kroonbladen.
petals

Nine obligatory boundaries in the reference occurred before a PP, which was used either attributively or non-attributively.[10] According to $\mathcal{N}\&\mathcal{V}$, however, no boundary can occur between a noun and an attributive PP, as this would break up an NP. In (8), for instance, the algorithm fails to predict a boundary between the noun *klimplant* and the attributive PP *met lange kantige en vertakte stengels*. The absence of a boundary after *klimplant* in combination with the presence of a boundary after *lange* – caused by list restructuring – would result in an odd kind of phrasing. The fact that a boundary is required here may again be explained by the length and/or complexity of the premodifier *borstelig behaarde* in combination with that of the postmodifier *met lange kantige en vertakte stengels*.

(8) Heggerank | is een borstelig behaarde klimplant #/∅ met
Bryony is a setaceously hairy climber with
lange | kantige | en vertakte stengels | . . .
long angular and branching stalks

A different problem occurs with non-attributive PP's. Since $\mathcal{N}\&\mathcal{V}$ only provides for restructuring after an NP or an S, it follows that only PP's preceded by one of these categories can begin a π. In Dutch, however, so called *extraposed* PP's are preceded by the verb(s) of the verb cluster. Consequently, $\mathcal{N}\&\mathcal{V}$ is unable to account for a boundary before an extraposed PP. An example of this is (9), for which the reference has an obligatory boundary between the verb *plaats* and the extraposed PP *met losbloemige bebladerde trossen van purperroze bloemen*.[11]

(9) Bloei vindt in de zomer plaats #/∅ met losbloemige
bloom does in the summer occur with polysepalous
bebladerde trossen van purperroze bloemen #
foliate bunches of purple-pink flowers

The third type of discrepancy is the context of two coordinated constituents. Example (10) shows coordinated NP's, and (11) two coordinated VP's. In both cases, the conjuncts are separated by an obligatory boundary in the reference, but $\mathcal{N}\&\mathcal{V}$ cannot account for such a boundary. In (10), a boundary would break up the dominating NP, whereas in (11), the boundary is after the verb,

[10] An *attributive* PP modifies a noun, whereas a *non-attributive* PP modifies a verbal or adjectival predicate.

[11] As a matter of fact, *plaats* is only the first part of the separable verb *plaatsvindt* 'takes place', lit. 'place-takes'. The final part, *vindt*, has moved to the second position of the root sentence, directly after the subject NP.

a position that the algorithm does not accommodate for.

(10) Ze bezitten afstaande kelkbladen #/∅ en vijf groenachtig
they possess protruding sepals and five greenish
witte netvormig geaderde kroonbladen. #
white clathrate veined petals

(11) Hij wordt wel | in tuinen gekweekt #/∅ en in
it is sometimes in gardens grown and into
muntsaus verwerkt. #
mint-sauce processed

The last category of errors concerns boundaries before *restrictive* relative clauses. Since such boundaries interrupt the dominating NP, $\mathcal{N}$&$\mathcal{V}$ fails to account for them. However, they occur repeatedly in the reference. An example of this is (12). In fact, the only case in which a boundary before a restrictive relative clause does *not* appear is when the clause is very short, as with the clause *die erop landen* in (13).

(12) ... # en dragen halverwege een bloeikolf #/∅ die schuin
and carry halfway a bloom-cob which slanted
afstaat en volgepakt is met onaangenaam geurende |
stands and packed is with unpleasant smelling
geelachtige | tweeslachtige bloemen. #
yellowish androgynous flowers

(13) insekten ∅ die erop landen | blijven steken # ...
insects that on-it land get stuck

4.4.2.3 Obligatory in the reference, but optional according to $\mathcal{N}$&$\mathcal{V}$ (#/|)

As can be seen from Table 4.3, the majority of the boundaries that are obligatory in the reference but merely optional according to $\mathcal{N}$&$\mathcal{V}$ occur between the items of a list. The general pattern appears to be that if the items are sufficiently complex, which may be defined as 'containing more than one lexical head', list restructuring becomes obligatory. Alternatively, the complexity of a conjunct may be defined in terms of the number of accents that it contains. Example (14) contains a list with moderately complex items, which are separated by obligatory boundaries in the reference. As matter of fact, there are sentences containing lists with much more complex items, even with items that are lists themselves.

Table 4.3: Locations that were obligatory boundaries in the reference, but optional boundaries according to $\mathcal{N}\&\mathcal{V}$.

Syntactic configuration:	*Number:*
between list items	27
before a non-attributive PP	9
between two conjuncts	
between two PP conjuncts	2
between two VP conjuncts	1
after a subject NP	3
Total:	42

(14) Aarmunt groeit op vochtige plaatsen #/ | langs
Mentha-Spicata grows in humid places along
wegen #/ | en op ruderale terreinen in grote delen van
roads and on ruderal grounds in large parts of
Europa # ...
Europe

Another class of problems are obligatory boundaries before non-attributive PP's. Although these PP's are often preceded by an NP, and consequently by an optional boundary, the reference frequently requires an obligatory boundary before them. For instance, $\mathcal{N}\&\mathcal{V}$ inserts an optional boundary after the NP *smaak en geurmiddel* in (15), while this ought to be an obligatory boundary in order to separate it from the long non-attributive PP *in snoepgoed, likeuren, medicijnen en tandpasta.*

(15) Deze olie | wordt gebruikt als smaak en geurmiddel #/ | in
this oil is used as flavour and aromatic-substance in
snoepgoed | likeuren | medicijnen | en tandpasta. #
candy liqueurs medicine and toothpaste

Yet another type of discrepancy can be found in contexts of the coordination of two constituents. In (16), for instance, the optional boundary between the two PP conjuncts should be an obligatory one according to the reference.

(16) De plant | groeit in venen en natte heide #/ | of op natte
the plant grows in peat and wet heath or on wet
rotsen ...
rocks

Table 4.4: Locations that had no boundaries in the reference, but were optional boundaries according to $\mathcal{N}$&$\mathcal{V}$.

Syntactic configuration:	*Number:*
after a subject NP	34
after a non-attributive PP	9
between two PP conjuncts	5
after a non-subject NP	1
before an S	1
Total:	50

Finally, there are a number of sentences in the corpus where a subject NP must be followed by an obligatory boundary, as in (17). Again, this may be attributed to the complexity of the constituent.

(17) De geelgroene bovenzijde van de bladschijf #/ | is bedekt
the yellow-green top of the leaf-slice is covered
door een weke | slijmige laag # ...
by a pappy slimy film

4.4.2.4 None in the reference, but optional according to $\mathcal{N}$&$\mathcal{V}$ ($\emptyset$/ |)

The bulk of the predicted optional boundaries that do not correspond to a boundary in the reference occur after a subject NP (see Table 4.4). Since a lot of subject NP's are rather short and uncomplicated – typically only a determiner and a noun – it makes sense that they are not followed by an optional boundary. According to $\mathcal{N}$&$\mathcal{V}$, however, a subject is not an internal argument, and therefore insensitive to the constraint that forbids separating an internal argument from its verb. As a matter of fact, $\mathcal{N}$&$\mathcal{V}$ would perform significantly better if this constraint were to extend to subjects as well. However, this would in turn introduce errors of another type, albeit less than the present 34. For instance, it would prevent the boundary after the subject NP in (17).

Some of the predicted optional boundaries are wrong because they create optional π's that lack a pitch accent. This is the case in (18), where an optional boundary after the subject NP *Deze olie* creates an unaccented π, which are generally considered to be unacceptable.[12]

(18) Deze olie $\emptyset$/ | wordt gebruikt als sm$\overset{*}{a}$ak en ge$\overset{*}{u}$rmiddel | ...
This oil is used as flavour and aromatic-substance

[12] There are a few exceptions to the rule; see Tone Copy in Gussenhoven (1992a, p. 96).

Other incorrectly predicted optional boundaries occur after NP's that function as complements in non-attributive PP's. Such a boundary occurs in (19) after the PP *in het wild.* For the same reasons, an optional boundary is predicted after the PP *om de pepermuntolie.* Since the verb *gekweekt* carries no pitch accent, the predicted boundary would give rise to an unaccented π.

(19) Hij groeit in het w$\overset{*}{\text{i}}$ld ∅/ | op v$\overset{*}{\text{o}}$chtige pl$\overset{*}{\text{a}}$atsen # en wordt
it grows in the wild in humid places and is
$\overset{*}{\text{o}}$ok om de peperm$\overset{*}{\text{u}}$ntolie ∅/ | gekweekt. #
also for the peppermint-oil grown

Finally, we find inappropriate optional boundaries between coordinated PP's. For example, (20) has incorrect boundaries between *in akkers* and *op braakliggende gronden*, as well as between *in wegbermen* and *langs rivieren.*

(20) ... # en is op veel plaatsen | ingeburgerd in akkers ∅/ | en
and is in many places established in fields and
op braakliggende grond | in wegbermen ∅/ | en langs
on fallow ground at roadsides and along
rivieren | ...
rivers

4.5 Intonational phrasing according to Caspers

4.5.1 $\mathcal{CAS}$ algorithm

Caspers performed a series of experiments about how Dutch speakers adapt their accent-lending and boundary-marking pitch movements during fast speech (Caspers 1994). As part of one of these experiments, she took three short texts (1,327 words, 52 sentences) read aloud by a professional speaker and transcribed all boundary-marking pitch movements. She also located all obligatory and optional π-boundaries as predicted by $\mathcal{N}\&\mathcal{V}$. Comparing the actual boundaries with the predicted boundaries, she found that 44% of the boundary-marking pitch movements occurred at places where no π-boundary was predicted (p. 37). Unpredicted boundaries were found before a PP (27%), before a restrictive relative clause (9%), after a conjunction (11%), between the items of a list of two (8%), and in indeterminate configurations (45%) (p. 186). In a follow-up experiment, she did the same for a short text (6 sentences) read aloud by six non-professional readers, although this time the beginning of a restrictive relative clause and the beginning of a PP counted as predicted optional boundaries as well. The number of unpredicted boundaries, at least for this text, dropped to 7% (p. 51). We will refer to this extension of $\mathcal{N}\&\mathcal{V}$ as $\mathcal{CAS}$. *π-formation* and *List restructuring* are identical, but *π-restructuring* is as follows.

(21) **π-restructuring** (optional):

a. Restructure after an NP if this does not interrupt another NP and does not separate and internal argument from its verb.

b. Restructure before an S.

c. Restructure before a PP.

Applying $\mathcal{CAS}$ to the sentence (1) gives (22). The difference with (21) consists of two extra optional boundaries: before the PP *in vele gebieden* and before the PP *in Nederland en Belgie.*

(22) De Reuzenbalsemien | is uit de Himalaya | naar Europa gebracht # en koloniseert hier waterkanten | in vele gebieden # ook plaatselijk | in Nederland en Belgie. #

4.5.2 $\mathcal{CAS}$ results

4.5.2.1 General

The overall results for $\mathcal{CAS}$ are presented in Table 4.5. The kappa value calculated for the matrix (without the optional boundaries) is identical to that for $\mathcal{N}\&\mathcal{V}$: the agreement between the prediction and the reference is very low and not significant at the 1%-level ($\kappa = 0.022$, $p = 0.045$). Recall that the part of the algorithm that decides on obligatory boundaries is the same for $\mathcal{N}\&\mathcal{V}$ and $\mathcal{CAS}$. Therefore, the predicted obligatory boundaries are identical to those for $\mathcal{N}\&\mathcal{V}$. However, in contrast to $\mathcal{N}\&\mathcal{V}$, $\mathcal{CAS}$ allows optional boundaries before PP's and restrictive relative clauses as well. As a consequence, it predicts more optional boundaries: sixteen obligatory boundaries that were not predicted at all by $\mathcal{N}\&\mathcal{V}$ are now predicted as optional by $\mathcal{CAS}$.

Contrary to what might be expected, the coverage of optional boundaries does not improve. From this, we can deduce two things. First, that there are no optional boundaries in the reference before a PP or a restrictive relative clause that were not already covered by $\mathcal{N}\&\mathcal{V}$. Second, there are still optional boundaries in the reference which are not located at the end of a NP, the beginning of an S, or the beginning of a PP. Thus, the extension of the $\mathcal{N}\&\mathcal{V}$ algorithm with respect to optional boundaries appears not to have the intended result.

In addition, the number of predicted optional boundaries where no boundary was allowed, which was 50 in $\mathcal{N}\&\mathcal{V}$, radically increases to 129 in $\mathcal{CAS}$.

In other words, a small improvement in the coverage of obligatory boundaries (realized as optional by $\mathcal{CAS}$ instead of not at all by $\mathcal{N}\&\mathcal{V}$) must be paid for with an excessive number of incorrect optional boundaries. This conclusion in fact agrees with the result of Caspers' experiments, which primarily addressed 'undershoot', i.e. actual boundaries in unpredicted locations; 'overshoot', i.e. predicted boundaries that are not actually realized, was of no direct concern. She observes, however, that one third of the predicted boundaries

Table 4.5: Confusion matrices for $\mathcal{CAS}$ algorithm. The top matrix contains the absolute number of confusions. The bottom matrix contains the percentage of confusions relative to the total number of boundaries in a column.

		Reference *Obl*	*Opt*	*Non*	*Tot*
$\mathcal{CAS}$	*Obl*	39	0	0	39
	Opt	58	34	129	221
	Non	12	8	575	595
	Tot	109	42	704	855

		Reference *Obl*	*Opt*	*Non*
$\mathcal{CAS}$	*Obl*	35.78	0	0
	Opt	53.21	80.95	18.32
	Non	11.01	19.05	81.68
	Tot	100	100	100

before a restrictive clause and almost three quarters of the boundaries before a PP did not receive a boundary-marking pitch movement (Caspers 1994, p. 51).

4.5.2.2 Obligatory in the reference, but none according to $\mathcal{CAS}$ ($\#/\emptyset$)

When we compare Table 4.6 to Table 4.2, we can conclude that the most important difference with $\mathcal{N}\&\mathcal{V}$ is that there are no missed obligatory boundaries before a PP or a restrictive relative clause. Moreover, the number of missed obligatory boundaries between AP's is halved. This is because some AP's are premodified by a PP, and consequently $\mathcal{CAS}$ places an optional boundary before them. Nevertheless, this result is not entirely satisfactory, for boundaries that $\mathcal{N}\&\mathcal{V}$ missed altogether are only predicted to be optional by $\mathcal{CAS}$, instead of obligatory.

4.5.2.3 Obligatory in the reference, but optional according to $\mathcal{CAS}$ ($\#/|$)

Errors of the type $\#/|$ are listed in Table 4.7. For several locations, the results are the same as for $\mathcal{N}\&\mathcal{V}$: between list items, between conjuncts and after a subject NP. However, the three obligatory boundaries before a restrictive relative clause, which $\mathcal{N}\&\mathcal{V}$ missed altogether (cf. Table 4.2), now appear as

Table 4.6: Locations that were obligatory boundaries in the reference, but no boundaries according to $\mathcal{CAS}$.

Syntactic configuration:	*Number:*
between AP's	4
between conjuncts	
between two NP conjuncts	6
between two VP conjuncts	2
Total:	12

Table 4.7: Locations that were obligatory boundaries in the reference, but optional boundaries according to $\mathcal{CAS}$.

Syntactic configuration:	*Number:*
between list items	27
before a PP	
before an attributive PP	6
before a non-attributive PP	12
before a PP modifying an A	4
between two conjuncts	
between two PP conjuncts	2
between two VP conjuncts	1
before a restrictive relative clause	3
after a subject NP	3
Total:	58

Table 4.8: Locations that had no boundaries in the reference, but were optional boundaries according to $\mathcal{CAS}$.

Syntactic configuration:	*Number:*
before a PP	
before an attributive PP	36
before a non-attributive PP	42
before a PP modifying an A	1
after a non-attributive PP	9
after a subject NP	34
between two PP conjuncts	5
after a non-subject NP	1
before an S	1
Total:	129

optional boundaries. The same goes for the six obligatory boundaries before an attributive PP, as well as for the four between AP's premodified by a PP.

4.5.2.4 None in the reference, but optional according to $\mathcal{CAS}$ ($\emptyset / \mid$)

The errors in Table 4.8 are to a large extent identical to the errors by $\mathcal{N}\&\mathcal{V}$ (cf. Table 4.4). As will be clear from a comparison of the two tables, the excessive number of incorrect optional boundaries is entirely due to those before a PP.

4.6 Intonational phrasing according to Gussenhoven

4.6.1 $\mathcal{GUS}$ algorithm

The previous two algorithms stem from prosodic phonology, which is primarily interested in prosodic phrasing at different levels and the way this affects the behaviour of segmental rules. It acknowledges the intonational phrase as one of these levels, but it tends to abstract away from intonational aspects such as pitch accent placement and the shape of the pitch contour, aspects of intonation that are addressed in the body of work that has recently become known as *intonational phonology.* Consequently, intonational phonology is in a position to criticise this abstraction by showing that the accent distribution and the pitch contour do affect intonational phrasing.

This is exactly the type of criticism that can be found in Gussenhoven (1992a, 1992a). Gussenhoven acknowledges that the morpho-syntactic structure is relevant to the derivation of intonational phrases, but shows that the pitch accent distribution is just as relevant. In fact, Gussenhoven takes the pitch accent distribution as the point of departure for intonational phrasing. In his view, every pitch accent needs to be realized in a certain part of the utterance. In autosegmental terms: the individual tones of a pitch accent need to be associated within a particular stretch of speech, which is called the pitch accent's *association domain* (AD).

AD's are derived by means of an algorithm that links them to both the accent distribution and the prosodic constituent structure of an utterance. Thus, the input to this phrasing algorithm consists of an accent distribution and a prosodic structure. As already mentioned in Section 4.3.4, this prosodic structure is somewhat different from the one assumed by $\mathcal{N}\&\mathcal{V}$ and $\mathcal{CAS}$. It does not conform to the *Strict Layer Hypothesis*, since recursion is allowed at the φ level.[13] Furthermore, its π's remain as derived by π-formation; there is no subsequent π-restructuring. Given these assumptions regarding the prosodic structure, the output of the algorithm consists of the locations of optional and obligatory AD boundaries.

It is important to keep apart the notions of π and AD. In Gussenhoven's view (cf. Section 2.3.7), π is still a part of the prosodic structure, but is no longer the domain of a single, coherent intonation contour; it is the AD which has taken over this property. The reason to preserve the π is that it constitutes a level of the prosodic structure, and is therefore relevant to the derivation of AD's and as the domain of application of certain segmental rules, just like the level of φ's is.[14]

As in the previous algorithms, phrasing proceeds in two steps. The first step, known as *AD derivation*, derives for every accent a corresponding unique boundary.

(23) **AD derivation**

For every accent, place an AD boundary at the end of the *highest* prosodic constituent that dominates it, but which does *not* dominate the next accent.

For the sentence in Figure 4.1, AD derivation results in the following AD's.

(24) De R$\overset{*}{\text{e}}$uzenbalsemien # is uit de Himal$\overset{*}{\text{a}}$ya # naar Eur$\overset{*}{\text{o}}$pa gebracht # en kolonis$\overset{*}{\text{e}}$ert # hier w$\overset{*}{\text{a}}$terkanten # in v$\overset{*}{\text{e}}$le # geb$\overset{*}{\text{i}}$eden # ook pl$\overset{*}{\text{a}}$atselijk # in N$\overset{*}{\text{e}}$derland # en B$\overset{*}{\text{e}}$lgie. #

The second step, called *AD restructuring*, joins one or more AD's into a

[13] See example (5) in Gussenhoven (1992a, p. 93). A strictly layered prosodic structure would in fact pose serious problems to AD-restructuring (Marsi, Coppen, Gussenhoven, and Rietveld 1997).

[14] Given this state of affairs, it is actually rather confusing to call an π an *intonational phrase*, and a different name would be more appropriate.

single one by deleting boundaries. This accounts for the observation that often several accents share a common AD. Restructuring is in part optional. Also, there is a constraint on the order of restructuring: precedence must be given to those boundaries that coincide with lower ranking prosodic constituents only (Gussenhoven and Rietveld 1992a, p. 288). This implies, for example, that φ-final AD boundaries must be deleted before π-final AD boundaries. Still, this formulation of restructuring is not specific enough to build an implementation on, because in principle every boundary can be deleted. Also, since there is no clear distinction between optional and obligatory boundaries, it difficult to compare the results to those of the two other algorithms. What is needed is some sort of principle to determine which boundaries will be deleted, and which boundaries will be turned into optional ones. For the purpose of evaluation, the following implementation of restructuring was used, which relies on both accent distribution and prosodic constituency.[15]

(25) **AD restructuring**

a. Delete all boundaries within a prosodic constituent that contains at most three accents.

b. Change all boundaries within a prosodic constituent that contains four accents into optional boundaries.

Restructuring the boundaries in (24) gives (26).

(26) De R*euzenbalsemien is uit de Himal*aya naar Eur*opa gebracht # en kolonis*eert | hier w*aterkanten in v*ele gebi*eden # ook pl*aatselijk in N*ederland en B*elgie. #

4.6.2 $\mathcal{GUS}$ results

4.6.2.1 General

The results for $\mathcal{GUS}$ are in Table 4.9. The kappa value calculated for the matrix (without the optional boundaries) indicates that the agreement between the prediction and the reference is rather high, and significant at the 1%-level ($\kappa = 0.876$, $p = 0.000$). In contrast to the other two algorithms, it predicts a considerable number of incorrect obligatory boundaries (33). On the other hand, it covers 86.24% of the obligatory boundaries in the reference. As for optional boundaries, $\mathcal{GUS}$ predicts only 40.48% of them. In exchange for this, the overshoot of optional boundaries is fairly limited in comparison to the other two algorithms.

Errors tend to fall in two categories. Recall that $\mathcal{GUS}$ requires a prosodic structure that is recursive at the φ-level. Accordingly, a number of the errors can be traced back to an inadequate prosodic structure. The majority of the

[15] It was found that using these numbers, i.e. three and four accents, gave the best results. However, there is no theoretical motivation for it.

Table 4.9: Confusion matrices for $\mathcal{GUS}$ algorithm. The top matrix contains the absolute number of confusions. The bottom matrix contains the percentage of confusions relative to the total number of boundaries in a column.

		Reference			
		Obl	*Opt*	*Non*	*Tot*
$\mathcal{GUS}$	*Obl*	94	12	21	127
	Opt	13	17	12	42
	Non	2	13	671	686
	Tot	109	42	704	855

		Reference		
		Obl	*Opt*	*Non*
$\mathcal{GUS}$	*Obl*	86.24	28.57	2.98
	Opt	11.93	40.48	1.70
	Non	1.83	30.95	95.31
	Tot	100	100	100

errors, however, is caused by an accent distribution that either incorrectly triggers restructuring or fails to trigger restructuring.

Some of the errors of $\mathcal{N}\&\mathcal{V}$ and $\mathcal{CAS}$ return in $\mathcal{GUS}$. For example, failure to predict obligatory boundaries between list items, and between two AP's. Other errors, like incorrectly predicted boundaries between conjuncts, do not recur. Also, due to the role of the accent distribution, creating intonational phrases without accents is avoided.

4.6.2.2 Obligatory in the reference, but none according to $\mathcal{GUS}$ ($\#/\emptyset$)

There were only two instances where $\mathcal{GUS}$ failed to predict an obligatory boundary (see Table 4.10). These two errors are in (27) and (28), where the b-parts show their prosodic structure. Recall that AD restructuring deletes all boundaries within a prosodic constituent that contains at most three accents. In (27), there should be an obligatory boundary after the subject NP *insekten die erop landen*, but since the dominating π contains only three accents, all boundaries inside it are deleted. In (28), the PP modifiers *in de bladoksels* and *aan verschillende planten* belong to a single φ that has just three accents. As a consequence, $\mathcal{GUS}$ is incapable of predicting the obligatory boundary that separates them. Interestingly, this boundary contributes to the disambiguation

Table 4.10: Locations that were obligatory boundaries in the reference, but no boundaries according to $\mathcal{GUS}$.

Syntactic configuration:	*Number:*
after subject NP	1
after non-attributive PP	1
Total:	2

of the sentence: it prevents an attributive interpretation of the final PP.

(27) a. Ins$\overset{*}{e}$kten die erop l$\overset{*}{a}$nden #/∅ blijven st$\overset{*}{e}$ken # ...
insects that on-it land get stuck

b. [$_\pi$
[$_\varphi$ Ins$\overset{*}{e}$kten]
[$_\varphi$ die erop l$\overset{*}{a}$nden] #/∅
[$_\varphi$ blijven st$\overset{*}{e}$ken]]

(28) a. De m$\overset{*}{a}$nnelijke en vr$\overset{*}{o}$uwelijke bl$\overset{*}{o}$emen # groeien in de
the male and female flowers grow in the

bl$\overset{*}{a}$doksels #/∅ aan versch$\overset{*}{i}$llende pl$\overset{*}{a}$nten. #
leaf-axils on different plants

b. [$_\pi$
[$_\varphi$
[$_\varphi$ de m$\overset{*}{a}$nnelijke]
[$_\varphi$ en vr$\overset{*}{o}$uwelijke bl$\overset{*}{o}$emen]]
[$_\varphi$ groeien]
[$_\varphi$
[$_\varphi$ in de bl$\overset{*}{a}$doksels] #/∅
[$_\varphi$ aan versch$\overset{*}{i}$llende pl$\overset{*}{a}$nten]]]

4.6.2.3 Obligatory in the reference, but optional according to $\mathcal{GUS}$ (#/ |)

Table 4.11 shows that a number of optional boundaries that should be obligatory occur in lists. For example, the list items in (29) constitute three separate φ's, which are in turn dominated by a higher φ. The latter contains four accents, which causes the boundaries separating the list items to become optional (cf. the principle of AD restructuring in Section 4.6.1). In other words,

Table 4.11: Locations that were obligatory boundaries in the reference, but optional boundaries according to $\mathcal{GUS}$.

Syntactic configuration:	*Number:*
between list items	6
between two AP's	4
between two VP conjuncts	2
before a non-attributive PP	1
Total:	13

whether the boundaries between list items are correctly predicted depends – – quite arbitrarily – on the number of accents in the list.

(29) a. ... op recht$\overset{*}{\text{o}}$pstaande | / ∅ $\overset{*}{\text{o}}$nbebladerde | / ∅ tot v$\overset{*}{\text{i}}$jftien
at erect leaf-less up-to fifteen
centimeter hoge st$\overset{*}{\text{e}}$ngels. #
centimetres high stalks

b. [$_\varphi$
[$_\varphi$ op recht$\overset{*}{\text{o}}$pstaande] | / ∅
[$_\varphi$ $\overset{*}{\text{o}}$nbebladerde] | / ∅
[$_\varphi$ tot v$\overset{*}{\text{i}}$jftien centimeter hoge st$\overset{*}{\text{e}}$ngels]]

For sequences of two AP's or VP's, the problem is essentially the same. Most of these do not contain enough accents to retain the obligatory boundary that, according to the reference, should separate their constituents.

4.6.2.4 None in the reference, but obligatory according to $\mathcal{GUS}$ (∅/#)

Table 4.12 lists all locations that were no boundaries in the reference, but obligatory boundaries according to $\mathcal{GUS}$. Recall that the dominance relations between φ's mirror the dominance relations between the syntactic constituents they are derived from (cf. Section 4.3.4). In particular, information about the level of attachment of an attributive PP remains encoded in the φ-structure. Consequently, a certain syntactic structure may lead to an inadequate φ-structure, which may in the case of $\mathcal{GUS}$ ultimately cause erroneous intonational boundaries. At least two of the incorrect boundaries before an attributive PP may in fact be traced back to the syntactic analysis.

To begin with, the boundary before the final PP in (30) – *in Europa* – would not appear if the prosodic constituents were as in (30-c) instead of (30-b),

Table 4.12: Locations that had no boundaries in the reference, but were obligatory boundaries according to $\mathcal{GUS}$.

Syntactic configuration:	*Number:*
before a PP	
before an attributive PP	6
before a non-attributive PP	6
after a subject NP	6
after a non-subject NP	2
after a verb	1
Total:	21

because all boundaries within a constituent containing at most three accents are deleted. This structure would arise if the PP were syntactically analysed as a modifier of *krijt en kalksteenbodems* rather than of *graslanden*. Thus, the question of whether this boundary is correct or not appears to depend on the syntactic analysis, and therefore ultimately on one's interpretation of the sentence.

(30) a. ... # in k*orte gr*aslanden # op kr*ijt en
in short grassland at chalk and
k*alksteenbodems ∅/# in Eur*opa. #
limestone-grounds in Europe

b. $[_\varphi$
$[_\varphi$ in k*orte gr*aslanden] #
$[_\varphi$ op kr*ijt en k*alksteenbodems] ∅/#
$[_\varphi$ in Eur*opa]]

c. $[_\varphi$
$[_\varphi$ in k*orte gr*aslanden] #
$[_\varphi$
$[_\varphi$ op kr*ijt en k*alksteenbodems] ∅/ ∅
$[_\varphi$ in Eur*opa]]]

The mirror image of this appears in (31). The boundary before the PP *van elliptische bladen* would be deleted with the prosodic analysis in (31-c) instead of that in (31-b), which would in turn require a syntactic analysis in which the PP *met een tot twintig centimeter hoge stevige bloeistengel* has scope over *een rozet van elliptische bladen*.

(31) a. … een roz$\overset{*}{\text{e}}$t $\emptyset$/# van ell$\overset{*}{\text{i}}$ptische bl$\overset{*}{\text{a}}$den # met een tot
a rosette of elliptical leaves with a up-to

tw$\overset{*}{\text{i}}$ntig centimeter h$\overset{*}{\text{o}}$ge | st$\overset{*}{\text{e}}$vige blo$\overset{*}{\text{e}}$istengel #
twenty centimeter high firm flower-stalk

b. [$_\varphi$
[$_\varphi$ een roz$\overset{*}{\text{e}}$t] $\emptyset$/#
[$_\varphi$ van ell$\overset{*}{\text{i}}$ptische bl$\overset{*}{\text{a}}$den] #
[$_\varphi$
[$_\varphi$ met een tot tw$\overset{*}{\text{i}}$ntig centimeter h$\overset{*}{\text{o}}$ge] |
[$_\varphi$ st$\overset{*}{\text{e}}$vige blo$\overset{*}{\text{e}}$istengel]]] #

c. [$_\varphi$
[$_\varphi$
[$_\varphi$ een roz$\overset{*}{\text{e}}$t] $\emptyset$/ $\emptyset$
[$_\varphi$ van ell$\overset{*}{\text{i}}$ptische bl$\overset{*}{\text{a}}$den]] #
[$_\varphi$
[$_\varphi$ met een tot tw$\overset{*}{\text{i}}$ntig centimeter h$\overset{*}{\text{o}}$ge] |
[$_\varphi$ st$\overset{*}{\text{e}}$vige blo$\overset{*}{\text{e}}$istengel]]] #

As for all other incorrect boundaries before an attributive PP, these are again a matter of too many accents, which trigger the deletion of the boundary.

Incorrect obligatory boundaries before non-attributive PP's invariably involve an extraposed PP, i.e. a PP after the verbal cluster containing the verbal or adjectival predicate. An example of this is (32). An alternative prosodic structure, where the φ's *is bedekt* and *door een weke slijmige laag* are joined into a higher φ, would solve the problem. Still, some of the extraposed PP's contain more than four accents, and will therefore be separated anyway.

(32) a. De g$\overset{*}{\text{e}}$elgroene b$\overset{*}{\text{o}}$venzijde van de bl$\overset{*}{\text{a}}$dschijf # is
the yellow-green top of the leaf is

bed$\overset{*}{\text{e}}$kt $\emptyset$/# door een w$\overset{*}{\text{e}}$ke sl$\overset{*}{\text{i}}$jmige la$\overset{*}{\text{a}}$g # …
covered by pappy slimy layer

b. [$_\pi$
[$_\varphi$
[$_\varphi$ De g$\overset{*}{\text{e}}$elgroene b$\overset{*}{\text{o}}$venzijde]
[$_\varphi$ van de bl$\overset{*}{\text{a}}$dschijf]] #
[$_\varphi$ is bed$\overset{*}{\text{e}}$kt] $\emptyset$/#
[$_\varphi$
[$_\varphi$ door een w$\overset{*}{\text{e}}$ke]
[$_\varphi$ sl$\overset{*}{\text{i}}$jmige la$\overset{*}{\text{a}}$g]]] #

Table 4.13: Locations that had no boundaries in the reference, but were optional boundaries according to $\mathcal{GUS}$.

Syntactic configuration:	*Number:*
before a PP	
before an attributive PP	5
before a non-attributive PP	2
after a subject NP	2
rest	3
Total:	21

Finally, incorrect obligatory boundaries occur after subject NP's because the following φ contains too many accents. This happens when one of the arguments is a list, as in (33).

(33) a. De bl$\overset{*}{\text{a}}$den ∅/# zijn gro$\overset{*}{\text{o}}$t | l$\overset{*}{\text{a}}$ngwerpig | l$\overset{*}{\text{a}}$nggesteeld |
the leaves are big elongated long-stalked
en get$\overset{*}{\text{a}}$nd. #
and serrated

b. $[_\pi$
$[_\varphi$ de bl$\overset{*}{\text{a}}$den] ∅/#
$[_\varphi$
$[_\varphi$ zijn gro$\overset{*}{\text{o}}$t] |
$[_\varphi$ l$\overset{*}{\text{a}}$ngwerpig] |
$[_\varphi$ l$\overset{*}{\text{a}}$nggesteeld] |
$[_\varphi$ en get$\overset{*}{\text{a}}$nd]]] #

4.6.2.5 None in the reference, but optional according to $\mathcal{GUS}$ (∅/ |)

These errors (see Table 4.13) are similar to those described in the previous section and their origin is essentially the same. The difference is that just three accents are involved, instead of four, so that the incorrectly predicted boundaries are merely optional instead of obligatory.

4.7 Intermediary discussion and summary

So far, three intonational phrasing algorithms have been evaluated. Although all of them perform reasonably well for the domain investigated, they show a number of systematic errors. $\mathcal{N}\&\mathcal{V}$ fails to predict boundaries, both obligatory and optional, because the syntactically defined set of possible boundary locations is too restricted. Furthermore, it predicts a considerable number of incorrect optional boundaries. $\mathcal{CAS}$ extends the set of possible optional boundary locations with the positions before a PP and before a non-restrictive S. Yet, this extension does not improve the coverage of optional boundaries, although it has the effect that a number of obligatory boundaries are now realized as optional instead of not at all. The price for this is a large number of incorrect optional boundaries.

In defence of Nespor & Vogel and Caspers, it should be noted that they never claimed that *all* their optional boundaries are actually valid. In their opinion, this depends on factors like speaking rate, length, style and prominence. For the same reason, some of their optional boundaries can in fact become obligatory. Such semantic and performance factors have not been formalised, and consequently could not be implemented. In general, intonational phrasing algorithms are claimed to predict *acceptable* boundaries, whereas our evaluation of them was directed toward predicting the *most acceptable* (preferred) boundaries. Still, the fact that the set of possible boundary locations is too restricted remains a serious problem. Neither $\mathcal{N}\&\mathcal{V}$ nor $\mathcal{CAS}$ can account for boundaries between AP's or conjuncts. Also, there is no provision to avoid empty intonational domains.

The main flaws of $\mathcal{GUS}$ are that it predicts too many obligatory boundaries and that it is not very good at predicting optional boundaries. The main reason for this is the lack of an adequate account of restructuring. The process depends on counting accents within a prosodic constituent, but the number of accents that triggers restructuring, i.e. three for obligatory restructuring and four for optional restructuring, is arbitrary. Yet, it is hard to think of a better account of restructuring, because the prosodic structure and the accent distribution are the only things available as input to $\mathcal{GUS}$. This makes it impossible to detect certain configurations – syntactic or otherwise – that require a particular intonational phrasing, as for lists, sequences of complex AP's, or coordinations. In other words, the prosodic structure is not *informative* enough for the assignment of intonational phrasing.

This finding is in line with the results of an experiment on AD restructuring in Marsi, Coppen, Gussenhoven, and Rietveld (1997), which showed that a prosodic structure without recursion at the φ-level[16] is not informative enough to account for the different preferences in AD-restructuring immediately before *attributive* and *non-attributive* PP's.

[16]That is, a prosodic structure that conforms to the Strict Layer Hypothesis, which states that a prosodic constituent of a certain rank cannot be dominated by prosodic constituent of the same rank, only by one of the next higher rank.

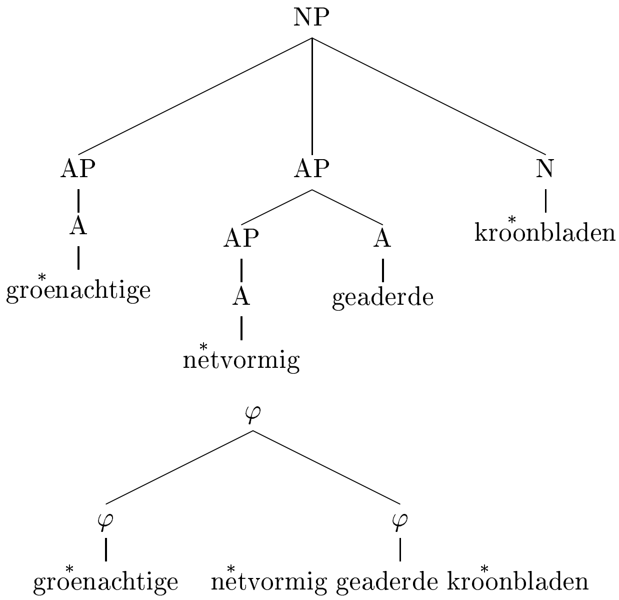

Figure 4.2: Syntactic structure and derived φ-structure for the phrase greenish clathrate veined petals.

With the current results, we can take this line of reasoning one step further, and argue that even a *recursive* prosodic structure is not informative enough for the purposes of intonational phrasing. In support of this position, consider the examples in Figure 4.2 and Figure 4.3: two different syntactic structures, which are nevertheless mapped to identical prosodic structures, i.e. with the same φ-structure and the same number of accents in each φ.[17] The former, however, requires an obligatory boundary between the two lower φ's according to the reference corpus, whereas the latter does not. Clearly, there is no way to explain this difference in intonational phrasing in terms of the prosodic structure and accent distribution only.

If we accept that our current prosodic structure is insufficient for the purpose of intonational phrasing, there are basically two ways to proceed. First, we can try to further extend our prosodic structure in such way that it incorporates the required information. This amounts to modifying the mapping rules so that even more semantic and syntactic details are preserved in the proso-

[17] The first example is derived from the following phrase in the reference corpus: ... # *en vijf groenachtig witte # netvormig geaderde kroonbladen*; lit. 'and five greenish white clathrate veined petals'. The second phrase does not occur literally in the corpus. However, of all the phrases in the corpus with an identical syntactic and prosodic structure, not one has an obligatory boundary between the two lower φ's.

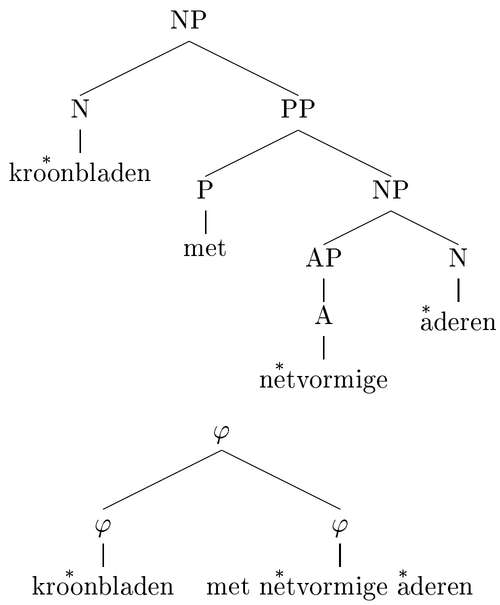

Figure 4.3: Syntactic structure and derived φ-structure for the phrase petals with clathrate veins.

dic structure. However, this raises a number of problems. To begin with, we would have to find a way to extend the prosodic structure, such as adding additional prosodic categories or allowing prosodic constituents to carry features. Moreover, these elaborations would require support by independent evidence, that is, we have to show that there are phonological rules that are sensitive to these extensions. Apart from these questions, this solution seems to defeat the whole purpose of the prosodic structure as an abstraction capturing only those semantic and syntactic aspects that are phonologically relevant. The prosodic structure would resemble the syntactic structure to such a degree that it would be hard to justify its existence. In conclusion, extending the prosodic structure is not an attractive solution.

The alternative way to proceed is to abandon the idea that intonational phrases are derived from the prosodic structure. Instead, we assume that they are derived directly from the syntactic structure. The rules for their derivation and restructuring have access to all the syntactic details that were shown to be relevant in the previous sections. At the same time, we can keep the mapping rules and the prosodic structure simple, in accordance with the Strict Layer Hypothesis as defended by Nespor & Vogel. As a drawback, we will probably lose the theoretical elegance of Gussenhoven's algorithm. However, arguments of theoretical economy have to give way to coverage of the empirical facts.

This second approach is clearly preferable to the first option outlined above. After all, since intonational phrasing – as one of the aspects of intonation – is inherently more closely related to *meaning* than purely *formal* phonological rules, it is a plausible assumption that it depends on input that contains a fair amount of meaning, i.e. fine-grained syntactic information.

So what we end up with is a combination of the best properties of the three algorithms. We agree with Nespor & Vogel and Caspers that intonational phrases are determined on the basis of the semantic and syntactic structure. However, the phrasing algorithm as proposed by Nespor & Vogel is in need of substantial modifications in order to increase its coverage and accuracy. Moreover, we preserve the two key insights of Gussenhoven, which may very well turn out to make a significant contribution to increasing the precision.

First, we reject the idea that a stretch of speech that corresponds to a single coherent intonation contour can be defined as a prosodic constituent with a fixed rank in the prosodic hierarchy (Section 2.3.7). Instead, we assume that it must be analysed as a separate entity of a different order. Following Gussenhoven, we may call it an *association domain*. Although the name of π, i.e. the prosodic category called *intonational* phrase, has thereby become somewhat misleading, we nevertheless preserve it as a prosodic category, because it is still relevant as the domain of application for certain phonological rules and for the overall temporal structure of the utterance (Gussenhoven and Rietveld 1992a). However, when there is no risk of confusion, we will use the terms *intonational phrase* and *association domain* indiscriminately. The same goes for *intonational phrasing* and *derivation of AD's*.

Second, we agree with Gussenhoven that the accent distribution plays a crucial role in intonational phrasing. In this way, phrases without accents can be avoided, and the number of accents can affect restructuring.

Finally, it is likely that intonational phrasing depends to some extent on speaking rate (Caspers 1994) and the length or weight of an intonational phrases (Marsi, Coppen, Gussenhoven, and Rietveld 1997) as well.

In conclusion, the results point in the direction of a syntax-driven algorithm for intonational phrasing, which is sensitive to a number of particular syntactic configurations, the accent distribution, and perhaps the speaking rate and weight of intonational phrases. An algorithm adhering to these principles will be presented in the next Section.

4.8 A new approach to intonational phrasing

4.8.1 $\mathcal{NEW}$ algorithm

This section proposes a new intonational phrasing algorithm called $\mathcal{NEW}$. It takes a syntactic tree and an accent distribution as input, and inserts intonational boundaries in the tree. It consists of eight rules for inserting intonational boundaries. Each rule describes a certain syntactic pattern plus conditions and the place to insert the boundary. There are dedicated rules for inserting a boundary in an apposition, a sentence, a verb phrase, a noun phrase, and a coordination (coordinated structure). An overview of all rules is given in Table 4.14. The rules apply in the order in which they are presented.[18] Each rule is applied obligatory and exhaustively before the next one is tried.

For all rules described below, it is assumed that intonational boundaries within coordinations are handled exclusively by the rule for this type of structures. This means, for instance, that the rules for phrasing in NP's do not match to a coordinated NP structure, as this would lead to unintended results. In addition, it is assumed that both the start and the end of an utterance constitute a trivial AD-boundary.

The rules are described and motivated with the help of examples. They are based on exemplary data, a considerable part of which is inspired by, or directly stems from the corpus of flower descriptions used in the above evaluation. The algorithm can thus to some extent be considered the result of manual modelling on the corpus of flower descriptions. It is therefore to be expected that a similar evaluation of $\mathcal{NEW}$ as described in Section 4.8.2, showed that it performs significantly better than two of the other algorithms. We will return to the question to what extent the algorithm is domain-independent Section 4.8.2.2.

4.8.1.1 Appositions

The class of appositional elements includes non-restrictive relative clauses, parenthetical expressions, tag questions, vocatives and expletives. There is no consensus about how appositions should be syntactically analysed. We will assume that they can take any syntactic category and can be attached anywhere in the syntactic tree, as long as they can be identified by the feature `appos`. With respect to the intonational phrasing of these elements, linguistic intuition appears to be unambiguous: appositional elements must be enclosed by intonational boundaries. This is corroborated by the results in Section 4.4.2, which clearly demonstrated that appositions were indeed always accompanied by intonational boundaries in the reference phrasing. We may explain this observation by the fact that appositional elements are typically signalled by lowering of the register, which by definition constitutes an intonational boundary.[19] Therefore, intonational phrasing for appositions is relatively simple, and is formalised in the following rule.

(34) **AD-to-Apposition**

For every apposition:
insert an intonational boundary before and after it.

[18] Only a few rules crucially dependent on their mutual ordering; these dependencies will be mentioned below. The rest of them is order independent in the sense that changing the order of the rules in the grammar will not affect the result.

[19] Although register lowering is not obligatory. For instance, tag questions can very well go without it.

Table 4.14: Rules of the $\mathcal{NEW}$ algorithm for AD-derivation.

Name:	*Definition:*
AD-to-Apposition	For every apposition: insert an intonational boundary before and after it.
AD-in-S-1	For every topicalised constituent: insert an intonational boundary after it, if (a) it has at least three accents, and (b) the subsequent AD has at least two accents.
AD-in-S-2	For every extraposed constituent: insert a boundary before it, if (a) it has at least two accents, and (b) the preceding AD has at least two accents.
AD-in-VP	For every daughter of a VP: insert a boundary after it, if (a) it has at least two accents, and (b) the subsequent AD has at least two accents.
AD-in-NP-1	For every premodifier in an NP: insert a AD-boundary after it, if (a) it carries at least one accent, and (b) its end coincides with a φ-boundary, and (c) the subsequent AD has at least one accent.
AD-in-NP-2	For every NP: insert an AD-boundary between its head and a postmodifier, if (a) the postmodifier has at least three accents, and (b) the preceding AD has at least one accent.
AD-in-NP-3	For every NP: insert an intonational phrase boundary between its postmodifiers, if (a) the preceding AD has at least one accent, and (b) the subsequent AD has at least one accent.
AD-in-coordination	For every pair of conjuncts in a complex coordination: insert an intonational phrase boundary between them, if (a) the preceding AD has at least one accent, and (b) the subsequent AD has at least one accent.

4.8.1.2 Sentences

At the sentence level, there are two structural positions at which an intonational phrase boundary may be inserted. The first is *after* a topicalised constituent, or – put differently – before the finite verb in a declarative main clause. The second is *before* an extraposed constituent, which is after the VP for the first extraposed constituent, and between two extraposed constituents for all others.

Boundary insertion at these positions appears to be primarily determined by the *length* or *weight* of the AD's that are created. For example, the boundary after the topicalised constituent *de mannelijke en vrouwelijke bloemen* in (35-a) is perfectly acceptable. However, if the topicalised constituent is shorter and has only one accent (35-b), or the remaining part of the sentence is shorter (35-c), inserting a boundary is far less appropriate, as this would results in relatively short AD's.

(35) a.

de	m$\overset{*}{\text{a}}$nnelijke	en	vr$\overset{*}{\text{ou}}$welijke	bl$\overset{*}{\text{oe}}$men	#	verschijnen
the	male	and	female	flowers		appear

d$\overset{*}{\text{oo}}$rgaans	aan	de	st$\overset{*}{\text{e}}$ngeltop
normally	at	the	stalk-top

b.

?de	bl$\overset{*}{\text{oe}}$men	#	verschijnen	d$\overset{*}{\text{oo}}$rgaans	aan	de	st$\overset{*}{\text{e}}$ngeltop
the	flowers		appear	normally	at	the	stalk-top

c.

?de	m$\overset{*}{\text{a}}$nnelijke	en	vr$\overset{*}{\text{ou}}$welijke	bl$\overset{*}{\text{oe}}$men	#	zijn	hetz$\overset{*}{\text{e}}$lfde
the	male	and	female	flowers		are	identical

It would appear, therefore, that both the weight of the AD *preceding* the boundary and the weight of the AD *following* the boundary have to be taken into account. The same type of *balancing* will return in most of the other rules described below. The main difficulty is to formalise the intuitive notion of an AD's *weight* by expressing it in terms of some unit of measure. We found that counting the number of pitch accents in a AD provides a reasonable approximation, and is relatively easy to implement.

Using this measure, the following rule inserts an AD boundary after a topicalised constituent.

(36) **AD-in-S-1**

For every topicalised constituent:
insert an intonational boundary after it, if

a. it has at least three accents, and
b. the subsequent AD has at least two accents.

With the *subsequent AD*, we mean the AD following the boundary, which comes into existence as a result of inserting this boundary. In (36-a), for example, the subsequent AD is *verschijnen doorgaans aan de stengeltop*. The condition that the topicalised constituent must have at least three accents is not derived from some theoretical principle; rather, it is a tendency observed in our corpus of flower descriptions. That is, choosing this value gives the best result when the algorithm is applied to the corpus. The same goes for the condition that the subsequent AD must have two or more accents.

Extraposed constituents show a behaviour that mirrors that of topicalised constituents in the sense that the decision to insert a boundary is affected by the weight of the extraposed constituent, as well as that of the constituent preceding it. In (37-a), for instance, both the VP and the extraposed PP contain two accents, and are therefore separated by a boundary. In contrast, the PP in (37-b) has only one accent, which prevents insertion of a boundary. The same goes for the VP in (37-c).

(37) a. hij komt in het *algemeen v*oor # in du*inen en z*anderige w*egbermen
he occurs in the general `prt` in dunes and sandy roadsides
'it normally occurs in dunes and sandy roadsides'

b. ?hij komt in het *algemeen v*oor # in du*inen
he occurs in the general `prt` in dunes

c. ?hij is gev*onden # in du*inen en z*anderige w*egbermen
he is found in dunes and sandy roadsides

On the basis of these observations, we propose the following rule.

(38) **AD-in-S-2**

For every extraposed constituent:
insert a boundary before it, if

a. it has at least two accents, and
b. the preceding AD has at least two accents.

Analogously to our use of the term *subsequent AD*, the *preceding AD* means the AD that is created to the left of the inserted boundary. In (37-a), the preceding AD is *hij komt in het algemeen voor*.

The rule for extraposed constituents is very similar to that for topicalised constituents. The difference is that the former requires the constituent to have at least three accents, whereas the latter requires at least two accents. The difference again stems from comparing the relevant instances in the corpus of flower descriptions.

In the implementation, the rules AD-S-1 and AD-S-2 are ordered. As a consequence, the former may destroy the context for the latter. In (39), for example, AD-in-S-1 inserts a boundary after the topicalised constituent, and therefore, AD-in-S-2 cannot insert a boundary before the extraposed constituent, as the preceding AD of AD-in-S-2 lacks the minimal number of accents.

(39) in Nederland, Duitsland, en Frankrijk # komt hij voor in
in Netherlands Germany and France occurs he `prt` in
duinen en zanderige wegbermen
dunes and sandy roadsides

However, switching around the rules gives acceptable results as well. With respect to (39), this means that there would only be a boundary before the extraposed constituent. As matter of fact, both boundaries occur when the number of accents in the intermediate part is increased (40).

(40) in Nederland, Duitsland, en Frankrijk # komt hij algemeen
in Netherlands Germany and France occurs he generally
voor # in duinen en zanderige wegbermen
`prt` in dunes and sandy roadsides

It is not uncommon to have more than one extraposed constituent in a sentence. AD-in-S-2 will insert a boundary between them provided that both carry at least two accents. This is true for virtually every extraposed constituent, since they are S's, or NP's or PP's that are 'heavy', which appears to be highly correlated with having multiple accents. The idea that the syntactic structure is affected by phonological considerations is also found in Zec and Inkelas (1990).

(41) hij komt algemeen voor # in Nederland, Duitsland, en
he occurs generally `prt` in Netherlands Germany and
Frankrijk # in duinen en zanderige wegbermen
France in dunes and sandy roadsides

Another interesting point is that the rules do not allow for an AD boundary after a fronted finite verb, i.e. a verb taking the second position in a declarative main clause or the first position in a polar question. Thus, an example as (42) is ruled out.

(42) *de mannelijke en vrouwelijke bloemen verschijnen #
the male and female flowers appear
doorgaans aan de stengeltop
normally at the stalk-top

When a boundary does follow the verb, this is in a sense a coincidence, since the boundary was actually inserted before an extraposed constituent. The boundary in (43), for instance, is inserted before the extraposed S.

(43) $(_S$ $(_{PP}$ uit rec$\overset{*}{\text{e}}$nte t$\overset{*}{\text{e}}$llingen) $(_V$ blijkt) $(_{VP}$) # $(_S$ dat de
from recent countings shows that the
plant in $\overset{*}{\text{a}}$antal ter$\overset{*}{\text{u}}$gloopt))
plant in number decreases

4.8.1.3 Verb Phrases

There are two general observations to make about intonational phrasing at the VP level. First, it is very unusual for an AD-boundary to occur before the verb cluster (the cluster at the end of the VP consisting of a main verb plus any verbal particles and auxiliaries). Second, a boundary may occur between phrasal daughters of the VP, depending on their weight. Again, the number of accents in a constituent appears to be a reasonably good indicator of its weight. Both observations are captured in the following rule.

(44) **AD-in-VP**

For every daughter of a VP:
insert a boundary after it, if

a. it has at least two accents, and
b. the subsequent AD has at least two accents.

Examples of boundaries due to AD-in-VP are given in (45). In (45-a), a boundary is inserted between the AdvP, which has two accents, and the PP, which has two accents as well. Example (45-b) is similar, except that the subsequent AD contains two daughters of the VP, each with only one accent. Example (45-c) shows that there is no boundary following the second PP *op veel plaatsen*, even though it contains two accents, because otherwise the subsequent AD would have just one accent. Finally, (45-d) has two boundaries due to AD-in-VP.

(45) a. die verandert $(_{VP}$ $(_{AdvP}$ tw$\overset{*}{\text{e}}$e maanden l$\overset{*}{\text{a}}$ter) # $(_{PP}$ in
which changes two months later in
een pl$\overset{*}{\text{u}}$izige b$\overset{*}{\text{o}}$l))
a fluffy ball

b. de plant komt $(_{VP}$ $(_{PP}$ in N$\overset{*}{\text{e}}$derland en B$\overset{*}{\text{e}}$lgie) #
the plant occurs in Netherlands and Belgium
$(_{AdvP}$ z$\overset{*}{\text{e}}$lden) v$\overset{*}{\text{o}}$or)
rarely `prt`

c. de plant komt $(_{VP}$ $(_{PP}$ in N$\overset{*}{\text{e}}$derland en B$\overset{*}{\text{e}}$lgie) #
the plant occurs in Netherlands and Belgium
$(_{PP}$ op v$\overset{*}{\text{e}}$el pl$\overset{*}{\text{a}}$atsen) v$\overset{*}{\text{o}}$or)
at many places `prt`

d. de plant komt ($_{VP}$ ($_{PP}$ in N$\overset{*}{\text{e}}$derland en B$\overset{*}{\text{e}}$lgie) #
the plant occurs in Netherlands and Belgium
($_{PP}$ op z$\overset{*}{\text{o}}$nnige pl$\overset{*}{\text{a}}$atsen) # ($_{AdvP}$ regelm$\overset{*}{\text{a}}$tig) v$\overset{*}{\text{o}}$or)
at sunny places regularly `prt`

All these examples comply with our first observation above that a boundary before the verb cluster is exceptional. This automatically follows from the second part of AD-in-VP. First, a verb cluster is usually followed by an AD-boundary, either because it is sentence-final, or because it is followed by an extraposed constituent. Second, the verb cluster will virtually always have at most one accented word, because auxiliaries are rarely accented. Taken together, this means that the condition in AD-in-VP that the subsequent AD must have at least two accents cannot be fulfilled. At the same time, this predicts that there will be two exceptions. The first occurs when the verb cluster does contain two accents, a type of accentuation that has strong pragmatic implications; see (46).

(46) dat een gr$\overset{*}{\text{o}}$otschalige versch$\overset{*}{\text{u}}$iving # z$\overset{*}{\text{o}}$u kunnen pl$\overset{*}{\text{a}}$atsvinden
that a large-scale shift may could take-place
is nog g$\overset{*}{\text{e}}$en bew$\overset{*}{\text{i}}$js
is yet no proof
'that a large-scale shift *might* occur is no proof yet'

The second exception occurs when the verb cluster is not closed by an AD-boundary, and the AD it is part of continues into an extraposed constituent. An example of this is (47-a). In contrast, insertion of a boundary is blocked in (47-b), because an intervening AD-boundary is present before the extraposed PP *als kleurstof*; a boundary which is due to the rule AD-in-S-2.

(47) a. de plant wordt in het z$\overset{*}{\text{u}}$iden van Eur$\overset{*}{\text{o}}$pa # gebr$\overset{*}{\text{u}}$ikt als
the plant is in the south of Europe used as
kl$\overset{*}{\text{e}}$urstof
colouring

b. de plant wordt in het z$\overset{*}{\text{u}}$iden van Eur$\overset{*}{\text{o}}$pa gebr$\overset{*}{\text{u}}$ikt # als
the plant is in the south of Europe used as
kl$\overset{*}{\text{e}}$urstof in v$\overset{*}{\text{o}}$edsel
colouring in food

To sum up so far, we have argued that intonational phrasing at the level of S and VP can be accounted for by just three rules, and that their ordering neatly accounts for the frequently observed balancing of AD's according to their weight.

4.8.1.4 Noun Phrases

There are two positions at which intonational boundaries in an NP may be observed: after a premodifier, i.e. after a daughter preceding the head, and before a postmodifier, i.e. before a daughter following the head.

With respect to the premodifiers, it appears to be the case that only AP's and NumP's are capable of triggering an intonational break, whereas other premodifiers, e.g. determiners and quantifiers, lack this ability (48-a). Moreover, a boundary can only occur between two subsequent premodifiers; never between the last premodifier and the head of the NP (48-b). By implication, a single premodifier will never trigger a boundary.

(48) a.

?dit	geldt	voor	($_{NP}$	($_{Q}$	$\overset{*}{\text{a}}$lle	)	#	bek$\overset{*}{\text{e}}$nde	leden	van de
this	goes	for			all			known	members	of the

N$\overset{*}{\text{a}}$chtschadefamilie)
Nightshade-family

b.

*dit	geldt	voor	($_{NP}$	$\overset{*}{\text{a}}$lle	($_{AP}$	r$\overset{*}{\text{e}}$chtstreekse	)	#	leden
this	goes	for		all		known			members

van de	N$\overset{*}{\text{a}}$chtschadefamilie)
of the	Nightshade-family

Both restrictions can be captured by the condition that the intonational phrase boundary must coincide with a phonological phrase boundary. Since there will never be a φ-boundary between the last premodifier and the head, there will never be an intonational boundary either. In the same way, determiners and quantifiers, which are function words, will never bring about a φ-boundary, and therefore never an AD-boundary either.

There is one complication: inserting a boundary after a premodifier may create an empty AD. Therefore, both the premodifier and the subsequent AD must have at least one accent. Taking all conditions together, we arrive at the following formulation.

(49) **AD-in-NP-1**

For every premodifier in an NP:
insert a AD-boundary after it, if

a. it carries at least one accent, and
b. its end coincides with a φ-boundary, and
c. the subsequent AD has at least one accent.

Postmodifiers – PP's and restrictive relative S's – may be preceded by an intonational boundary if sufficiently heavy. By inspecting the corpus of flower descriptions, we found that if the head of the NP is followed by a postmodifier that has at least three accents, head and modifier are separated by an AD-boundary.[20] As with the previous rule, creating an empty AD must be avoided.

[20] Of course, we found exceptions to this generalisation, but this value gives the highest

This can be accomplished by checking that the preceding AD has at least one accent.

(50) **AD-in-NP-2**

For every NP:
insert an AD-boundary between its head and a postmodifier, if

a. the postmodifier has at least three accents, and
b. the preceding AD has at least two accent.

Two examples of boundaries due to (50) are given in (51).

```
(51) a.  de  plant vormt (NP fl*inke       (N kol*onies ) # (PP in
         the plant forms      substantial     colonies            in

         b*ossen w*eilanden en  d*uinen ) )
         forest  meadows   and dunes

     b.  (NP een e*enjarige (N pl*ant ) # (S die    ho*oguit
             a   annual        plant          which maximally

         tw*intig centimeter ho*og wordt   ) )
         twenty   centimetre high  becomes
```

So far, only the *first* postmodifier has been addressed. However, an NP can have more than one postmodifier; see (52). The fact that the postmodifying PP's can be switched around without altering the meaning of the full NP suggests that we have indeed two independent constituents, and that it is not the case that one of them is embedded in the other.

```
(52) (NP een (N roz*et   ) (PP van ell*iptische bl*aden ) (PP met  een
         a      rosette        of  elliptical   leaves         with a

     tot   tw*intig centimeter h*oge st*evige blo*eistengel ) )
     up-to twenty   centimetre high  firm     flower-stalk
```

With regard to intonational phrasing, the best solution appears to be to always insert an intonational boundary between subsequent postmodifiers. Again, care must be taken not to create AD's without accents.

(53) **AD-in-NP-3**

For every NP:
insert an intonational phrase boundary between its postmodifiers, if

a. the preceding AD has at least one accent, and
b. the subsequent AD has at least one accent.

number of correctly predicted boundaries when applying the algorithm to the corpus.

These three rules suffice to account for intonational phrasing in NP's. Notice that – as with the rules for phrasing in S's and VP's – they are concerned with intonational boundaries between *direct* daughters of the NP only. These daughters – AP's, NP's, PP's, S's – may contain intonational boundaries themselves, which are due to other phrasing rules. This includes a recursive application of the AD-in-NP rules, as in the case of PP-embedding.

4.8.1.5 Coordinations

There appear to be a number of ways to phrase a coordination. The first is to map the whole coordination onto a single AD and to apply the tonal operation *Downstep* (Section 2.3.6), which results in one type of 'list intonation' (54-a). We may hypothesise that – at least in the case of coordinations – downstep facilitates parsing of the syntactic structure in the sense that each downstepped accent serves to signal a separate conjunct. This would explain the observation that downstep is less likely to be used for a coordination in which one of the prefinal conjuncts has more than one pitch accent (54-b). For some reason, the last conjunct of a coordination forms an exception to the rule (here *het noorden van Luxemburg)*, as it may contain multiple accents even when downstep is applied (54-c).

(54) a. op n*atte gr*ond ! # in ($_{NP}$ ($_{NP}$ N*ederland) ($_{NP}$
at wet ground in the-Netherlands
B*elgie) en ($_{NP}$ L*uxemburg))
Belgium and Luxembourg

b. ?op n*atte gr*ond ! # in ($_{NP}$ ($_{NP}$ het z*uiden van
at wet ground in the south of
N*ederland) ($_{NP}$ h*eel B*elgie) en ($_{NP}$ het no*orden
the-Netherlands whole Belgium and the north
van L*uxemburg))
of Luxembourg

c. op n*atte gr*ond ! # in ($_{NP}$ ($_{NP}$ N*ederland) ($_{NP}$
at wet ground in the-Netherlands
B*elgie) en ($_{NP}$ het no*orden van L*uxemburg))
Belgium and the north of Luxembourg

We will refer to a coordination in which none of the prefinal conjuncts carries more than one pitch accent as a *simple coordination.* No phrasing rule is required for simple coordinations; they remain free of internal AD-boundaries. Insertion of the Downstep operator before a simple coordination will be discussed in Section 7.4. The outer boundaries at the start and at the end of a coordination are taken care of by the other phrasing rules, depending on the context. For instance, the boundary before the coordination in (54) is due to

AD-in-NP-2.[21] In contrast, a coordination occurring directly before the verb cluster will not be closed by an AD-boundary.

The second approach to phrasing a coordination is to map each conjunct onto an AD of its own. This is appropriate when at least one of the prefinal conjunct has more than one accent. We will refer to these as *complex coordinations*. For this purpose, we introduce the following rule.

(55) **AD-in-coordination**

For every pair of conjuncts in a complex coordination:
insert an intonational phrase boundary between them, if

a. the preceding AD has at least one accent, and
b. the subsequent AD has at least one accent.

In theory, it is possible that one of the conjuncts is unaccented. Therefore, the rule checks that the created preceding and subsequent AD's are not empty. In practice, however, it is hard to find examples of this type. No examples of this type occurred in the corpus of flower descriptions. As for a simple coordination, the boundaries surrounding a complex coordination are due to independent phrasing rules.

Interestingly, the weight of a constituent, which is assumed to affect phrasing in most of the other syntactic contexts discussed above, does not play a part in the phrasing of coordinations. Instead, the decisive factor appears to be whether at least one of the prefinal conjuncts contains more than one accent.

With a view to the treatment of pitch accent placement in numerals in Section 7.3.5, we will finish this section with showing that AD-in-coordination is also responsible for intonational phrasing in complex numerals. In Section 6.2.3.4, we will argue that complex numerals have to be syntactically analysed as coordinated structures, even though the conjunction *en* 'and' can usually be omitted. An example of this is (56), of which the syntactic structure is given in Figure 4.4.

(56) honderd vier en vijftig duizend zeven honderd twee en dertig
hundred four and fifty thousand seven hundred two and thirty
kasten
cupboards
'154,732 cupboards'

There are two alternative accentuations of (56) which appear to be equally acceptable: with an accent on *duizend* and a following AD-boundary (57-a), or without either an accent or a following AD-boundary (57-b).

[21] Strictly speaking, the boundary is not inserted before the coordinated NP structure itself, but rather before the PP that dominates it. These are details of the implementation that will be ignored here.

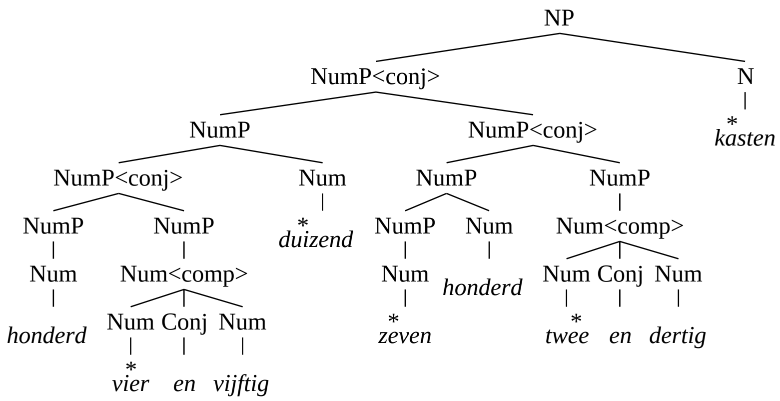

Figure 4.4: Syntactic structure of example (56).

(57) a. [$_{\text{AD}}$ honderd v$\overset{*}{\text{i}}$eren vijftig d$\overset{*}{\text{u}}$izend]
[$_{\text{AD}}$ z$\overset{*}{\text{e}}$ven honderd tw$\overset{*}{\text{e}}$e en dertig k$\overset{*}{\text{a}}$sten]

b. [$_{\text{AD}}$ honderd v$\overset{*}{\text{i}}$eren vijftig duizend z$\overset{*}{\text{e}}$ven honderd tw$\overset{*}{\text{e}}$e en dertig k$\overset{*}{\text{a}}$sten]

It will be clear that no AD boundaries are predicted in the lower coordinations, because the number of accents in their penultimate conjuncts is insufficient. However, an accent on *duizend* provides the penultimate conjunct of the highest coordination with the required number of accents to trigger AD-in-coordination. Thus, the rule AD-in-coordination, in combination with the syntactic analysis of complex numerals as involving coordination, correctly predicts dependency relations between accent and boundary as observed in (57). We will return to this issue in Section 7.3.5

4.8.2 $\mathcal{NEW}$ results

4.8.2.1 General

Earlier this chapter, the intonational phrasing algorithms $\mathcal{N}\&\mathcal{V}$, $\mathcal{CAS}$, and $\mathcal{GUS}$ were evaluated by applying them to a corpus of flower descriptions and comparing their output to a reference phrasing produced by human experts. Applying the same evaluation procedure to $\mathcal{NEW}$, we obtain the result summarised in Table 4.15, which shows the difference between the prediction and the reference in the form of a confusion matrix. The kappa value calculated for the matrix (without the optional boundaries) indicates that there is a high agreement between the prediction and the reference that is significant at the 1%-level

Table 4.15: Confusion matrices for $\mathcal{NEW}$ algorithm. The top matrix contains the absolute number of confusions. The bottom matrix contains the percentage of confusions relative to the total number of boundaries in a column.

Reference

$\mathcal{NEW}$		*Obl*	*Opt*	*Non*	*Tot*
	Obl	98	10	9	117
	Opt	0	0	0	0
	Non	11	32	695	738
	Tot	109	42	704	855

Reference

$\mathcal{NEW}$		*Obl*	*Opt*	*Non*
	Obl	89.91	23.81	1.28
	Opt	0	0	0
	Non	10.09	76.19	98.72
	Tot	100	100	100

($\kappa = 0.923$, $p = 0.000$). The second row contains only zeros, because $\mathcal{NEW}$ – in contrast to the other algorithms – does not assign optional boundaries. The second column can be ignored, because at locations where an optional boundary occurs in the reference, both a boundary and no boundary are acceptable, that is, the algorithm cannot be wrong. $\mathcal{NEW}$ assigns 9 erroneous boundaries and omits 11 boundaries. Another way to look at this is that of all 855 boundary decisions, only 20 (2.34%) are wrong. Considering the length and complexity of the sentences in the corpus, this is a highly accurate result.

It is interesting to compare these results with those obtained with the previously tested algorithms. Of these, $\mathcal{GUS}$ performed best; see Table 4.9. However, a comparison is not straightforward, because $\mathcal{GUS}$ assigns optional boundaries, whereas $\mathcal{NEW}$ does not. Therefore, the optional boundaries have to be transformed to either obligatory boundaries or non-boundaries. There are three options:

1. Promote all optional boundaries to obligatory boundaries, i.e. add the second row to the first in Table 4.9.

2. Reduce all optional boundaries to non-boundaries, i.e. add the second row to the third one in Table 4.9.

3. Divide all optional boundaries equally over the obligatory boundaries and the non-boundaries, i.e. add half of the second row to the first one, and add the other half to the third one in Table 4.9. This amounts to a random choice for each optional boundary to either realize or omit it.

The results for each of these three options are given in Table 4.16. The differences among the three options with respect to the total number of errors turn out to be negligible. The general picture that emerges is that on the one hand $\mathcal{NEW}$ misses a few more boundaries than $\mathcal{GUS}$, but on the other hand assigns considerably fewer incorrect boundaries. Still, the total number of errors for $\mathcal{NEW}$ is 20, whereas that for $\mathcal{GUS}$ is at least 35.

A comparison of $\mathcal{NEW}$ with the other two algorithms ($\mathcal{N}\&\mathcal{V}$ and $\mathcal{CAS}$) is less interesting, because the differences are considerably bigger, regardless of the option we choose for counting optional boundaries. For instance, the total number of errors for $\mathcal{N}\&\mathcal{V}$ is minimally 70 (8.19%).

This interpretation of the differences between the algorithms was confirmed by a statistical analysis. For every algorithm, we calculated its kappa value (discounting optional boundaries) for each of the ten texts. An analysis of variance (repeated measurements) carried out on these kappa values showed that the factor algorithm was significant at $F_{3,27} = 24.140$, $p < 0.01$ (Huynh-Feldt corrected). A subsequent post-hoc comparison procedure yielded two homogeneous subsets: $\{\mathcal{N}\&\mathcal{V}, \mathcal{CAS}\}$, and $\{\mathcal{GUS}, \mathcal{NEW}\}$ ($p < 0.05$).

4.8.2.2 Discussion

$\mathcal{NEW}$ is based on exemplary data, a considerable part of which is inspired by, or directly stems from the corpus of flower descriptions used in the above evaluation. The algorithm is in part the result of repeating the cycle of testing and improving until the point that – in our opinion – no further improvement was feasible. Although it is possible to improve the results for $\mathcal{NEW}$ with additional phrasing rules, these rules would then account for just one or two

Table 4.16: $\mathcal{NEW}$ compared to $\mathcal{GUS}$ with respect to undershoot (#/$\emptyset$) and overshoot ($\emptyset$/#) errors for obligatory boundaries. Percentages are calculated with respect to the total number of boundary decisions (855).

Algorithm:	*Undershoot:*		*Overshoot:*		*Total:*	
$\mathcal{NEW}$	11	(1.29%)	9	(1.05%)	20	(2.34%)
$\mathcal{GUS}_1$	2	(0.23%)	33	(3.86%)	35	(4.09%)
$\mathcal{GUS}_2$	15	(1.75%)	21	(2.46%)	36	(4.21%)
$\mathcal{GUS}_3$	8.5	(0.99%)	27	(3.16%)	35.5	(4.15%)

extra correct boundaries. Clearly, this sort of evidence is insufficient to justify such extensions. $\mathcal{NEW}$ is thus the result of balancing economy of description on the one hand and accurate prediction on the other.

Consequently, we cannot present an analysis of systematic errors produced by $\mathcal{NEW}$. After all, if we were able to find patterns or classes of errors, we would seize the opportunity to improve the algorithm. Instead, the errors are all rather individual cases that might be attributed to an inadequate measure of AD weight, the possibility of alternative syntactic interpretations, special pragmatic contexts, explicit attempts to avoid ambiguity, or an incidental case of inconsistent judgements of the panel that produced the reference phrasing. Whatever our speculations may be, these are cases for further research with the help of more data (see Appendix B).

Since the algorithm is to some extent the result of manual modelling on the corpus of flower descriptions, it is not surprising that evaluation of $\mathcal{NEW}$ against this corpus showed that it performs better than at least two of the other algorithms. The question is to what extent a corpus of ten short texts, all of the same genre, is representative of written texts in general. In other words: does $\mathcal{NEW}$ perform equally well when applied to other types of written text? Ultimately, this question can only be answered by repeatedly evaluating the algorithm against many different corpora, something which is beyond the scope of the current work. Large speech corpora becoming available nowadays are probably not going to be of any help with respect to generating the most acceptable intonation. To begin with, they are usually not prosodically annotated. In addition, they are typically designed for wide coverage, containing speech from competent ('good') as well as less competent ('bad') speakers. Apart from the problem that the necessary semantic and syntactic information are probably not available, extracting regularities from the corpus will inevitably lead to a model of the intonational behaviour of an average speaker, and the generation of *largely acceptable* rather than *good* intonation.

Although the evaluation of $\mathcal{NEW}$ has not been repeated with another corpus, the algorithm has been incorporated in our SLG system (to be described in the next chapters). In addition to flower descriptions (WFD), this system generates travel descriptions (PTD). Results on an evaluation of the generated travel descriptions will be reported in Chapter 8. Part of this concerns the intonational boundaries as predicted by $\mathcal{NEW}$.

4.9 Summary

The goal of this chapter was to obtain an accurate intonational phrasing algorithm, i.e. one that can account for the intonational boundaries that are most acceptable according to intonational experts. In addition, we intended to find out which types of information are required as input. Although a considerable amount of linguistic theory regarding intonational phrasing is available in the literature, most work is based on a small set of crucial observations and has neither been tested against larger sets of data nor systematically compared with competing approaches. In the first part of this chapter, therefore, the algorithms by Nespor & Vogel ($\mathcal{N}$&$\mathcal{V}$), Capers ($\mathcal{CAS}$), and Gussenhoven ($\mathcal{GUS}$) were systematically compared using a substantial set of data. We used a corpus-based method in which the intonational boundaries predicted by the algorithm were compared with those prescribed by a panel of experts, where the text material consisted of flower descriptions. A detailed error analysis revealed systematic errors in each of the algorithms. It was concluded that a better algorithm would have to be syntax-driven, and sensitive to a number of particular syntactic configurations, the accent distribution, and the weight of intonational phrases. Accordingly, a new algorithm ($\mathcal{NEW}$) was proposed consisting of eight rules for inserting intonational boundaries in particular syntactic configurations. It correctly predicted 97.66% of the boundary decisions in the corpus of flower descriptions, which was shown to be significantly better than the predictions by $\mathcal{N}$&$\mathcal{V}$ and $\mathcal{CAS}$. An additional evaluation of the $\mathcal{NEW}$ algorithm, as incorporated in our SLG system, will be described in Chapter 8.