CONTENTS

Towards Radically Incremental Parsing of Natural Language

Wolfgang Menzel
Hamburg University

Abstract

Human sentence processing proceeds in a left-to-right incremental manner and is able to assign partial structural interpretations even if all the lexical items involved are not yet known. To mimic such a behaviour in a computational system is particularly difficult for head final languages where the rich information about the valency of the verb comes last. In such cases the Argument Dependency Model facilitates early hypotheses about the thematic status of a constituent by establishing direct argument-to-argument dependencies independent of the verb. An implementation of the model within the framework of Weighted Constraint Dependency Grammar (WCDG) is presented and evaluated in a pseudo-incremental processing mode. Due to its non-monotonic nature, the system can replicate the predictions of the Argument Dependency Model about the dynamic nature of revision processes during sentence comprehension with a very high degree of reliability.

1 Introduction

One of the most obvious characteristics of human language understanding is its incremental nature. Humans are processing sentences in a left-to-right manner making choices as early as possible always risking the necessity of later revisions, which might become necessary to accommodate the subsequent input information. Indeed, the vast majority of psycholinguistic insights into the mechanisms of the human language faculty relate to the time course of language processing. In such a situation, any attempt to compare the performance of the human model with that of artificial systems beyond a simple introspective account requires algorithmic solutions with comparable processing properties. While natural language parsing technology has made significant progress with respect to its coverage and reliability, the lack of a truly incremental processing mode remains one of its most serious drawbacks.

The motivation to investigate prerequisites and possibilities for incremental processing of natural language utterances, however, reaches far beyond a cognitive perspective. Left-to-right processing is also an essential feature for all kinds of online processing tasks, which are particularly relevant in speech processing scenarios, like dialogue or dictation. Here, the

speaking time becomes a precious resource, which needs to be used in order to maintain fluency and provide for a natural man-machine interaction.

This paper investigates possibilities for incremental processing in the context of a broad coverage parsing model implemented within the framework of Weighted Constraint Dependency Grammar (WCDG) (Schröder 2002, Foth 2007).

2 Incremental sentence processing

Incremental processing of an input stream is an inherently recursive process: After a certain portion of the input has been analysed, it is extended by additional elements (an increment) and subjected to the very same processing component. This process becomes non-monotonic as soon as the necessity arises to revise (parts of) the already computed output information according to the extended input. Two different approaches can be distinguished in such a case: a cautious one, which maintains all the different output alternatives until enough information is available to take a safe decision, and an eager one, which adopts the most plausible interpretation as soon as possible and does so even at the risk that later revisions become necessary. Obviously, human sentence processing takes the latter approach. This does not come as a surprise, because eager processing has a number of important cognitive advantages:

1. Taking early decisions provides for a focussed, and, hence, more efficient allocation of cognitive resources.
2. Having a single preferred reading available early enough allows the hearer to derive strong expectations about upcoming observations. Maintaining a variety of pending hypotheses will inevitably result in very many different expectations, which are not only expensive to be matched to the actual input, but due to their diversity have little predictive power.
3. Having a unique interpretation available rather early enables the hearer to react quickly to the incoming information, an important prerequisite for effective communication strategies.

Even among the eager approaches different temporal processing schemes can be distinguished. Usually, a piece of the structural description is built as soon as all of its constituting parts are available. Unfortunately, such a schema necessarily involves some kind of delay. While in the case of a phrase structure grammar all daughter nodes of a rule to be applied need to be available, for a dependency model at least the two word forms to be connected by a dependency relation must be already known. Even if incomplete partial structures are hypothesized on a sub-rule level, e.g., by means of chart parsing techniques, a final decision on their appropriateness

can only be taken after the complete constituent has been established.

This situation is particularly problematic in head-final languages since crucial information about valences and valence requirements, hinging upon the head (namely the verb), comes last. Therefore, reliable hypotheses on the distribution of functional roles can only be established at a very late point in time. This again is in stark contrast to human sentence processing where hypotheses, for instance about the thematic status of a constituent, are established very early and independent of the verb and its lexical information. They are revised later if verb-specific information, like government or agreement requirements, has to be accomodated.

Recently, large coverage models of natural language syntax have become available, which are able to process sentences with almost no failure and a fairly high accuracy. One of them, MaltParser (Nivre et al. 2007), even adopts a strictly incremental processing scheme. Here, a stack serves as an intermediate storage device able to hold unattached word forms until an appropriate attachment point becomes available. For each incoming word form a support vector machine operating on features of the parse history deterministically decides which one of the following operations is applied:

- shift the incoming word form on the stack, i.e., wait
- attach the word form on the stack to the incoming one and reduce the stack
- attach the incoming word form to the one on the stack and shift it on the stack
- reduce the stack

Although the decision is fully deterministic, i.e., no revisions are ever considered, the parser achieves almost state-of-the-art performance for a number of languages (Nivre 2007). Still, it suffers from the delay which is incurred by shifting input elements onto the stack temporarily. Changing this behaviour to an earlier commitment would require an additional source of information about role attachments independent of the verb and a non-deterministic mechanism for revising decisions if necessary.

A framework that does just this is the Argument Dependency Model (ADM). It establishes a hierarchy of direct complement-to-complement relationships which are fully independent of the valence requirements of individual verbs. This approach, therefore, not only makes verb-independent information available for disambiguation, but also facilitates an early assignment of a functional role to a certain part of the sentence even if the governing verb is not yet known. The model has been successfully used to explain EEG-data obtained from humans while listening to German subordinate clauses (Bornkessel 2002). To investigate whether the ADM might also help to facilitate early decisions and revision processes in a computational parser we integrated it into an existing broad coverage grammar for

German and evaluated it on the very same utterances that have been used in the psycholinguistic experiments of Bornkessel (2002).

3 The Argument Dependency Model

Psycholinguistic evidence gives rise to the expectation that the degree of meaning derived from a given sentence during incremental comprehension is a function of the morphological informativeness of the sentential arguments (Schlesewsky & Bornkessel 2004). Accordingly, the Argument Dependency Model (Bornkessel et al. 2005), which is based on the language comprehension model of Friederici (2002), has been designed as a dual-pathway architecture, consisting of two different processing components: a syntactic and a thematic one. Which of the two pathways is chosen depends on the morphological case marking borne by the incoming arguments: The thematic pathway is activated by an unambiguously case marked argument, otherwise the syntactic pathway is chosen. The choice of a pathway does not mean that the other pathway is cut off but that the chosen pathway leads the comprehension.

Thematic information provides a general conceptual specification of the relations between the arguments of a sentence and between the arguments and the verb. Hence, the thematic pathway generalizes over a number of verbs by being based on so called proto-roles. These proto-roles have been introduced as Proto-Agent, Proto-Patient, and Proto-Recipient by for instance Dowty (1991) and Primus (1999). Proto-roles as defined by Primus are dependent on each other and, thus, can be arranged according to a thematic hierarchy:

$$\text{Proto-Agent} <_{\Theta} \text{Proto-Recipient} <_{\Theta} \text{Proto-Patient}$$

The ideal Proto-Agent is nominative and in control (animate). It may not depend on any other argument ($-$DEP). In contrast, a Proto-Patient ($+$DEP) always depends on either a Proto-Agent or a Proto-Recipient. The Proto-Recipient, finally, can be either depending, dependent, or both at the same time ($\pm$DEP).

This thematic hierarchy can be mapped to a complementing case hierarchy according to a many-to-many relationship: Nominative is an ideal proto-agent since it has full control.

Ich zerbrach die Vase. (*I broke the vase.*)

Dative can be proto-agent, but has no full control

Mir zerbrach die Vase.

**Mir zerbrach die Vase absichtlich.*

whereas accusative can never be proto-agent

**Mich zerbrach die Vase.*

Dependency relationships between the constituents of a sentence are established by means of a number of principles that either operate on the thematic pathway or on the syntactic one:

Principle	**Thematic pathway**	**Syntactic pathway**
ECONOMY	The first argument is assigned the status [−dep] if at all possible.	An argument is associated with the grammatical function compatible with the least number of syntactic dependencies.
DEPENDENCY	For any two arguments A and B, either A must hierarchically dominate B in terms of thematic status, or B must dominate A.	For any two arguments A and B, either A must c-command B, or B must c-command A.
DISTINCTNESS	For relations consisting of ≥ 2 arguments, each argument must be maximally distinct from every other argument in terms of thematic status.	For any two arguments A and B, either A must asymmetrically c-command B. or B must asymmetrically c-command A.

These principles can be refined for individual languages, e.g., for German

MAPPING: if no verb-specific information contradicts, nominative (+ANIMATE) receives −DEP
BLOCKING: accusative is not compatible with −DEP

They are complemented by the usual verb specific requirements (VERB-FIT), e.g.,

thematic: object-experiencer verbs require that the dative-marked argument must dominate the nominative-marked one
syntactic: subject-predicate agreement must be obeyed

Finally, a general preference for argument role attachment over adjunct attachments is assumed.

Most of these principles are non-monotonic ones. They assign default values (e.g., ECONOMY) or require that an optimum between different alternatives be determined (e.g., DISTINCTNESS). It is this non-monotonicity that provides an explanation for the re-interpretation processes that can be observed with human sentence processing.

In the following example the role assignment solely relies on the thematic

pathway. After the first constituent (*der Junge/the boy*) is encountered, it is assigned the topmost rank in the hierarchy (−DEP). Consequently, the second one receives +DEP according to the distinctness principle. No reanalysis occurs since the assignment is compatible with the case of the second argument and the specific requirements of the verb.

..., dass	*der Junge*	*den Großvater*	*besucht.*	
	−DEP			ECONOMY
		+DEP		DISTINCTNESS
				VERB-FIT
... that	*the boy*$_{\text{NOM}}$	*the grandfather*$_{\text{ACC}}$	*visits.*	

In the second example, a reanalysis is necessary on the second constituent, because its nominative case ultimately requires −DEP and therefore the first constituent needs to be reanalysed as +DEP. The object-experiencer verb finally triggers yet another re-interpretation since it forces the dative to dominate the nominative.

..., dass	*dem Jungen*	*der Film*	*gefällt.*	
	−DEP			ECONOMY
		−DEP		MAPPING
	+DEP			DISTINCTNESS
	−DEP			VERB-FIT
		+DEP		DISTINCTNESS
..., that	*the boy*$_{\text{DAT}}$	*the movie*$_{\text{NOM}}$	*likes.*	

A similar kind of reanalysis is also necessary with ternary verbs, where the maximally distinct assignment on the dative NP needs to be corrected as soon as the accusative becomes available. If (in the case of ambiguous case assignments) the syntactic pathway is activated, reanalysis can only be triggered by the agreement requirements of the verb.

4 Incremental processing with CDG

A preliminary investigation using a Weighted Constraint Dependency Grammar (WCDG) has shown that in principle the model can be used in an incremental processing mode (Foth et al. 2000a). This study, however, has been carried out using a fairly restricted model for a specific type of utterances (namely simplified utterances from an appointment negotiation domain). Moreover, similar to MaltParser (Nivre et al. 2007), this system did not implement a truly incremental processing scheme but had to wait instead until a suitable attachment point became available somewhere in the right

context. A partial parsing scheme was used to deal with the incomplete nature of a sentence during incremental analysis. It relaxes the constraints on possible top nodes of a dependency tree, so that arbitrary categories can take this role and the tree breaks into fragments if no better attachment was available. This in some cases led to a need for unintuitive weight assignments in the grammar.

It has been shown that the parser was indeed robust enough to deal with many instances of sentence prefixes, i.e., incomplete utterances. Moreover, its ability to supply information about constraint violations in addition to structural hypotheses turned out to be an extremely valuable feature in the incremental case: Constraint violations for an incomplete utterance reflect expectations for the not yet observed input and may therefore guide further processing.

Trying to adopt this approach in a broad coverage grammar, however, seems to be not that easy. At least two aspects contribute to these difficulties:

- The parser enforces attachment decisions for all the word forms in the utterance. This might easily lead to meaningless attachments as long as the sentence is still unfolding, because a suitable attachment point might not yet be available (see the upper part of Figure 1). The low plausibility of this structure is also reflected in its score which is more than an order of magnitude smaller than that of the second alternative below. To avoid such unnatural intermediate solutions a processing mechanism is required which explicitly reflects the (possibly) incomplete nature of the incoming utterance.
- Considering the usually very large variety of possible continuations for a sentence prefix, many constraints have to be relaxed in order to tolerate missing information. Such a neutralization of syntactic constraints might be possible in a restricted domain, but has disastrous effects with a broad coverage grammar. Here, the restrictive potential of grammar constraints is already fairly low because a large number of alternatives has to be accommodated. Somehow, this loss of constraining information needs to be compensated for.

In principle, WCDG is highly compatible with the requirements of an eager processing mode:

- Irrespective of the input being partial or complete, WCDG is able to determine the optimal structural interpretation.
- Dependency structures in general are robust against fragmentation: Any part of a dependency tree is a valid dependency tree again. WCDG inherits this property.
- By evaluating individual dependency relations (or pairs of them) the scope of constraints is fine-grained enough to support local decisions.

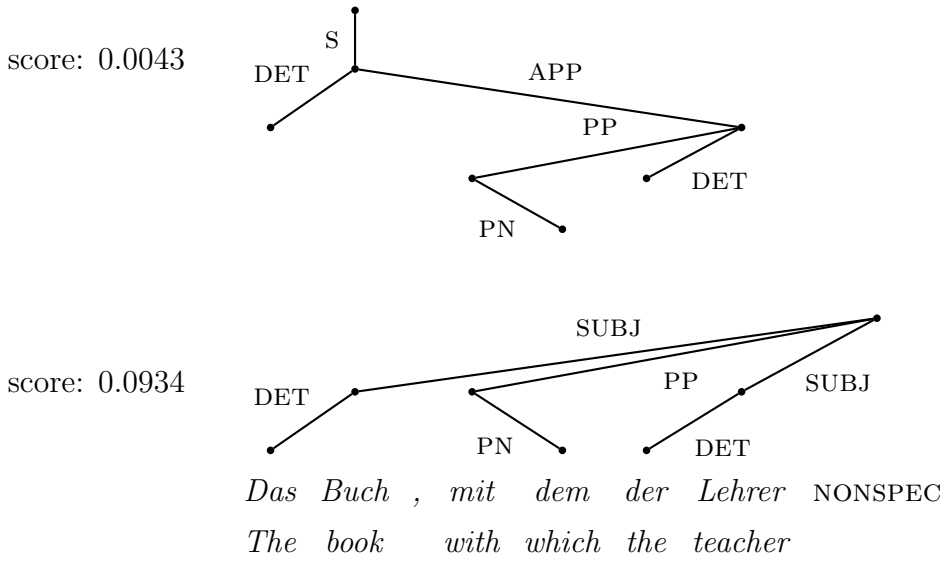

Fig. 1: *Structural interpretation of an incomplete sentence with and without* NONSPEC*-modeling*

> That is, the attachment of a right complement can be established without waiting for possible other complements not yet observed.

From a procedural point of view, algorithms of different degrees sophistication can be devised for incremental processing. They differ with respect to the amount of information that is passed between two subsequent calls on an incrementally extended input stream. The most simple though least interesting solution is prefix parsing, where for each new increment, the complete input sequence available so far is subjected to the parser again. No information about the results of a preceding call is maintained or passed to subsequent processing steps. To really reap the benefits of incremental processing, however, more information about the hypothesis space of the preceding cycle needs to be made available to the subsequent one.

Foth et al. (2000a) compared a series of increasingly permissive structural heuristics and found a tradeoff between speedup and quality. Note, however, that this comparison has been made using a small scale grammar and a combinatorial search procedure, which builds dependency trees by successively adding more edges. Such an approach turned out to be infeasible if applied to large scale grammars as used in the experiments reported here. For them a repair approach has been more successful, which sucessively transforms parts of a structure in order to remove the most severe constraint violations (Foth et al. 2000b). Such a transformational approach comes with an additional advantage since it lends itself for a completely

different kind of informational coupling between subsequent incremental cycles: Instead of (or in addition to) *restricting* the remaining search space, the currently available optimal structure for a sentence prefix can be used to *initialize* or even *guide* the transformation procedure in the subsequent step.

Using the optimal structure of the current step as a starting point for the next one was the approach taken in the experiments reported here. Since so far it does not restrict the transformation possibilities, only a fairly weak coupling between subsequent incremental cycles is provided. Therefore, this type of coupling is referred to as pseudo-incremental processing in the following. No impressive results in terms of its temporal characteristics could be expected. It is meant, however, as a first attempt to study possibilities of incremental processing with a large scale WCDG under the conditions of a truly eager processing regime.

In order to facilitate eager processing, the basic solution procedure of WCDG had to be extended as well to be able to consider also the not yet observed right context of the currently available sentence prefix. The ideal solution would consist in providing a set of virtual nodes which could be filled later with the lexical information of the incoming word forms. Two problems, however, have to be solved:

1. How many of these additional nodes are actually necessary?
2. Are there sufficient conditions under which two nodes can be identified as referring to the same input word form, and thus combined?

Because there are no straightforward heuristics to answer these questions, we adopted a compromise: A single node, called NONSPEC (Daum 2004) is used to approximately model the unseen right context and all word forms requiring a modifiee to the right are attached to it (see the lower part of Figure 1). Its approximate nature results from the fact that it merely serves as a projection plane for arbitrary right-bound attachment expectations and it remains unclear whether this node actually represents one or several different word forms. Moreover, the NONSPEC-node has no lexical properties and none can be projected onto it. Thus, its purpose consists only in avoiding meaningless attachments and a further propagation of their consequences across the already seen sentence prefix.

The NONSPEC node competes with other nodes in the parsing problem in the ususal manner: Attachments are established, scored and possibly modified if they turn out to be the source of a major constraint violation. From the perspective of an external observer only the consequences of this competition become visible as a kind of reanalysis effect, since the optimum might switch to an alternative structure as the scores of the individual dependency edges develop. Internally, it is just an accumulation of constraint violations at a certain point in the structure. i.e., a reranking of alternative

interpretations. Of course, the final optimal structure of a sentence should not include a NONSPEC node. This, however, is not directly enforced by the parser.

5 Experiments

5.1 *The grammar*

To facilitate experiments with an incremental processing mode, the original grammar as used e.g., in (Foth 2007, Foth & Menzel 2006) was

- modified to deal with NONSPEC and
- extended by an implementation of the Argument Dependency Model.

Since NONSPEC neither represents the identity of a word form nor does it contain any lexical information, all constraints of the grammar have to be adapted accordingly in order to avoid unwarranted conclusions or access to non-existing information. This, in particular, requires to prevent all constraints on right-bound dependencies from accessing lexical information at the NONSPEC-node.

While in Section 3 the Argument Dependency Model has been introduced by means of a feature assignment mechanism, its implementation in WCDG is based on true argument-to-argument relationships which directly represent the corresponding dependencies as postulated by the model. For this purpose, an additional level ARG has been introduced. It is used to build a chain of argument dependencies which strictly obeys the hierarchical ordering conditions of the DEPENDENCY principle by placing arguments higher up in the chain if they exhibit fewer dependencies than others. Constraints have been added to model the different principles of the Argument Dependency Model. Since the defeasible constraints of WCDG are fundamentally non-monotonic in nature, in most cases there is a one-to-one mapping between principles and constraints. For more details see (Menzel & Dalinghaus 2008).

Note that these modifications have been the only ones made to the original grammar.

5.2 *Data*

To facilitate a direct comparison with the available findings about the psychological adequacy of the Argument Dependency Model we used the same set of sentences, which have been compiled as test stimuli for the experiments with human subjects (Bornkessel 2002).

The corpus comprises sentences according to a uniform sentence pattern consisting of a verb final target subclause preceded by its matrix clause ⟨MC⟩ (e.g., *Gestern wurde gesagt/It was said yesterday*) and followed by a second subclause ⟨SC⟩ (e.g., *obwohl das nicht wahr ist/although this isn't true*).

⟨MC⟩ *dass* *die Winzer* *dem Betrüger* *abraten,* ⟨SC⟩
NP1 ↑ NP2 ↑ V ↑

that the winegrowers the swindler dissuade,

The relevant points in time, where decisions about the thematic status of constituents have to be taken or possibly revised are marked by arrows.

The target subclause is modified into 16 different test conditions, along the following four parameters

- constituent order: subject before object vs. object before subject

  ... dass Christian Professorinnen abrät
  vs. *... dass Christian Professorinnen abraten*

  ... that Christian$_{\mathrm{SG}}$ *dissuades*$_{\mathrm{SG}}$ *Professors*$_{\mathrm{PL}}$.
  vs. *... that Professors*$_{\mathrm{PL}}$ *dissuade*$_{\mathrm{PL}}$ *Christian*$_{\mathrm{SG}}$.
- verb type: active (*abraten/dissuade*) vs. object-experiencer (*gefallen/like*)[1]
- case marking: ambiguous (Hans/Hans, Bäuerinnen/peasants, ...) vs. unambiguous (dem Betrüger/the cheater)
- morphological variation: singular before plural vs. plural before singular (with a possible inflectional adaptation of the verb form to ensure subject-verb agreement).

For each of these 16 conditions 80 sentences have been generated yielding a corpus consisting of 1280 sentences, which can be used to systematically study different cases of reanalysis processes. Basically, two different kinds have to be considered here: syntactic and thematic revisions, where a syntactic reanalysis can only appear in case the first argument carries an ambiguous case marking.

If in any of these conditions a reanalysis becomes necessary it occurs during the integration of the verb information. A more fine grained investigation of the time course of revision processes can be conducted using the unambiguous conditions, where an early revision (at the second NP) can be distinguished from a late one at the verb (c.f. Figure 3).

5.3 *Results*

Four different research questions have guided the experimental work carried out in this study (Menzel & Dalinghaus 2008):

- Do the parsing decisions on incomplete sentences comply with the predictions of the Argument Dependency Model, i.e., can the reanalysis effects be observed in the intermediate parser output?

[1] Note that only dative objects are considered.

- How sensitive is the output quality to a modification of fundamental modelling assumptions?
- Does the pseudo-incremental parsing scheme provide a significant improvement over mere prefix parsing?
- Do the different reanalysis processes have an impact on the resource requirements of the parser?

To study these issues in more detail the test sentences were presented in three incremental steps to the parser: The sentence up to the point after the first NP, the second NP and the verb. The final subclause which was needed to factor out sentence final integration effects in the psycholinguistic experiments has not been considered here. The non-incremental mode was used as a baseline for comparison.

Using the predictions of the model as a gold standard a quality-centered evaluation can be conducted in a rather straightforward manner. At a first glance, the global quality of parsing decisions seems very high. With a completely correct structure for whole sentences (including the matrix clause and the final subclause) on both levels (syntax and thematic) in 97.3% of the test sentences it exceeds the quality level of current state-of-the-art dependency parsers on general text corpora by far. Given, however, the very regular structure of the test data this result is not really a surprising one. With few exceptions (e.g., the ambiguity at the verb *gefallen/like* between its object-experiencer interpretation and the participle form of *fallen/fall*) no serious parsing problems occur.

A closer inspection reveals that intermediate results at the critical points in the sentence are far less reliable (c.f. Figure 2). The treatment of the second NP seems particularly error prone. Nevertheless, for the majority of test sentences the interpretation-switching behaviour predicted by the Argument Dependency Model can be replicated by the WCDG parser during left-to-right incremental parsing.

The vast majority of error cases (93%) are confusions of an argument with either a genitive modifier (GMOD), an apposition (APP), or an ethical dative (ETH). They have been caused by a mismatch between fundamental assumptions of the Argument Dependency Model and the original WCDG grammar: While in the Argument Dependency Model argument attachments take priority over all other attachments, the WCDG grammar attaches a noun phrase as an argument only if this is licensed by the subcategorization pattern of the verb. This mismatch can explain the comparatively low accuracy at the second NP, since at that position the verb information is not yet available. As a consequence, the parser runs into systematic problems in cases with a nominative-genitive or genitive-dative syncretism, since here an alternative non-argument interpretation is readily available. The problem is particularly severe as the test data set does not contain any sentences with non-argument NPs, making the parser fail on all the relevant instances.

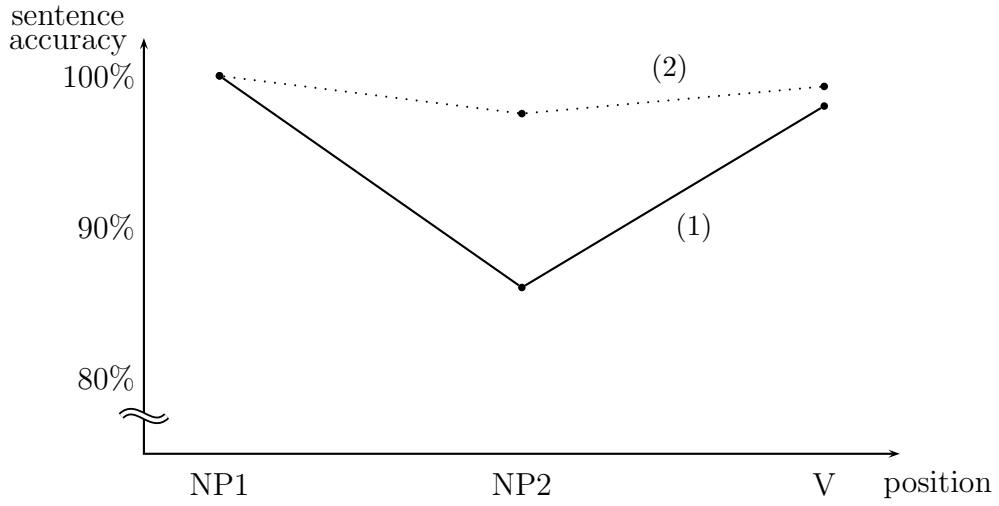

Fig. 2: *Parsing quality at the critical points with the original grammar (1) and the reduced preference for non-argument attachments (2)*

Therefore, an additional experiment has been conducted to estimate the potential for improvement. It shows that a very simple modification of the grammar (penalizing the non-argument attachment alternatives) removes the problem almost completely. The dotted line in Figure 2 shows the corresponding result for deprecating the labels GMOD and ETH.[2] Similar results can also be expected for arguments wrongly attached as appositions to the preceding noun. This finding confirms the assumption that there is potential to further improve the parsing results in accordance with the predictions of the Argument Dependency Model. Considered in isolation, however, the result is not very meaningful, since a modified preference of the grammar might of course negatively affect its performance on a general purpose corpus. Whether this is the case and if so, whether an acceptable balance can be found remains as a goal for further investigations.

In general, it is not clear how the parsing quality of an incremental model can be evaluated on a standard treebank at all. Thus, transferring the current implementation of the Argument Dependency Model to the case of unrestricted text as used in standard evaluations poses yet another challenge, since no gold standard for parsing unrestricted text in an incremental manner is available so far. To demonstrate the appropriateness of the model derived for the controlled conditions of the stimulus sentences also under

[2] Note that the measures for sentence accuracy in Figure 2 cannot be compared directly, since the second experiment has been carried out on only a subset of 160 sentences (10 for each condition).

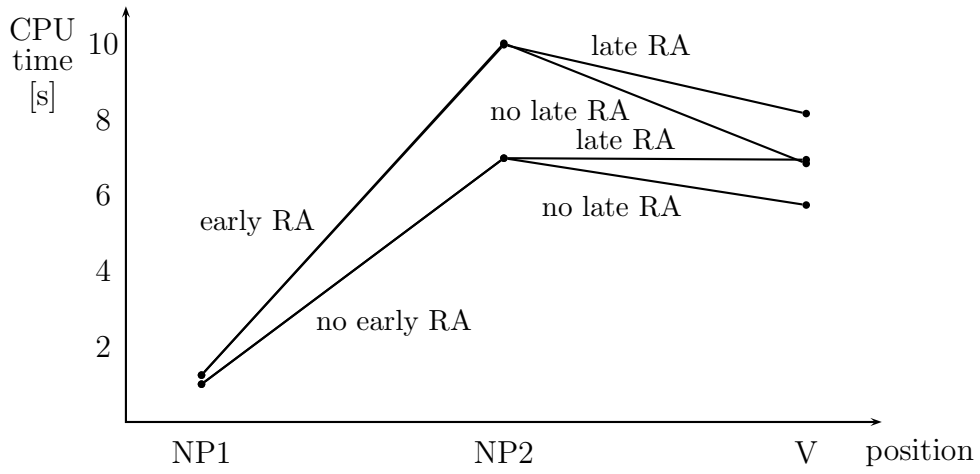

Fig. 3: *Computational effort at the critical points for sentences with the singular NP preceding the plural*

the open world conditions of general text data, it would be necessary to make assumptions about the intermediate interpretations produced by humans during online comprehension. The Argument Dependency Model is only of limited use in that situation.

With respect to the temporal behaviour of the parser, less convincing results have been found so far. In particular no noticeable speed-up compared to pure prefix parsing was measured. Obviously, the pseudo-incremental parsing scheme which was used in the experiments provides too weak a coupling between subsequent incremental processing steps. For some of the experimental conditions, reanalysis effects have clearly been reflected in the run-time requirements of the parser. Figure 3 shows the CPU time consumed for processing the three increments up to the critical points in the sentence. In these cases, the need for reanalysis (RA) is systematically correlated to higher resource requirements, and the higher computational effort is triggered exactly at the point in the sentence at which the evidence becomes available. Unfortunately, this behaviour does not extend to the complementary conditions in which the plural noun phrase precedes the singular.

6 Conclusions

A model that captures dependencies between the arguments of a common head has been implemented using a parsing formalism based on weighted constraints. Modeling such dependencies is motivated by the need to compensate for the loss of disambiguating information in a partially available

sentence, thus allowing well-informed early decisions about the type of attachment even in cases where the attachment point has not yet been observed. Compared to previous approaches a considerably more radical kind of incremental sentence processing can be achieved.

It has been shown that weighted constraints are an appropriate means to model grammar in a way that mirrors human preferences about the interpretation of incomplete sentences and that produces, as a consequence of their application, the non-monotonic syntactic and thematic re-interpretations of the sentence as it unfolds. Applying a suitably extended broad-coverage parser to the very same test data as have been used in psycholinguistic experiments, it was possible to show that this pattern of early commitment can be computationally replicated with a very high degree of reliability.

Acknowledgements. This research was partly funded by the Deutsche Forschungsgemeinschaft under grant ME 1472/6-1. I am grateful to Klaus Dalinghaus who implemented the grammar extensions and conducted the experiments described above.

REFERENCES

Bornkessel, Ina. 2002. "The Argument Dependency Model: A Neurocognitive Approach to Incremental Interpretation". Dissertation, Universität Potsdam, Germany.

Bornkessel, Ina, Stefan Zysset, Angela D. Friederici, D. Yves von Cramon & Matthias Schlesewsky. 2005. "Who Did What to Whom? The Neural Basis of Argument Hierarchies during Language Comprehension". *NeuroImage* 26:1.221-223.

Daum, Michael. 2004. "Dynamic Dependency Parsing". *Proceedings of ACL Workshop on Incremental Parsing*, 67-73. Barcelona, Spain.

Dowty, David. 1991. "Thematic Proto-Roles and Argument Selection". *Language* 67:3.547-619.

Foth, Kilian A. 2007. *Hybrid Methods of Natural Language Analysis.* Aachen: Shaker Verlag.

Foth, Kilian & Wolfgang Menzel. 2006. "Hybrid Parsing: Using Probabilistic Models as Predictors for a Symbolic Parser". *Proceedings of the 21st International Conference on Computational Linguistics* (*Coling-ACL 2006*), 321-328. Sydney, Australia.

Foth, Kilian A., Wolfgang Menzel, Horia F. Pop & Ingo Schröder. 2000a. "An Experiment in Incremental Parsing using Weighted Constraints". *Proceedings of the 18th International Conference on Computational Linguistics* (*Coling 2000*), 1026-1030. Saarbrücken, Germany.

Foth, Kilian A., Wolfgang Menzel & Ingo Schröder. 2000b. "A Transformation-based Parsing Technique with Anytime Properties". *Proceedings of the 4th*

International Workshop on Parsing Technologies (*IWPT 2000*), 89-100. Trento, Italy.

Friederici, Angela D. 2002. "Towards a Neural Basis of Auditory Sentence Processing". *Trends in Cognitive Science* 6:2.78-84.

Menzel, Wolfgang & Klaus Dalinghaus. 2008. "An Implementation of the Argument Dependency Model". Report FBI-HH-B-284, Department Informatik, Universität Hamburg.

Nivre, Joakim. 2007. "Incremental Non-Projective Dependency Parsing", *Proceedings of the Human Language Technologies: The Annual Conference of the North American Chapter of the Association for Computational Linguistics*, (*NAACL-HLT-2007*), 396-403. Rochester, NY.

Nivre, Joakim, Johan Hall, Jens Nilsson, Atanas Chanev, Gülşen Eryiğit, Sandra Kübler, Svetoslav Marinov & Erwin Marsi. 2007. "MaltParser: A Language-independent System for Data-driven Dependency Parsing". *Natural Language Engineering* 13:2.95-135.

Primus, Beatrice. 1999. *Cases and Thematic Roles: Ergative, Accusative and Active*. Tübingen: Max Niemeyer Verlag.

Schlesewsky, Matthias & Ina Bornkessel. 2004. "On Incremental Interpretation: Degrees of Meaning Accessed during Sentence Comprehension". *Lingua* 114:9-10.1213-1234.

Schröder, Ingo. 2002. Natural Language Parsing with Graded Constraints. Dissertation, Universität Hamburg, Fachbereich Informatik.

Index of Subjects and Terms