

From Elementary Discourse Units to Complex Ones

Holger Schauer

Computational Linguistics Division

Freiburg University

D-79085 Freiburg, Germany

schauer@coling.uni-freiburg.de

Abstract

Coherence relations have usually been taken to link clauses and larger units. After arguing that some phrases can be seen as discourse units, a computational account for such phrases is presented that integrates surface-based criteria with inferential ones. This approach can be generalized to treat intra-sentential cue-phrases. Since cue-phrases are not always present, referential relations between nominal expressions are additionally used to derive a text's discourse structure.

1 Introduction

It is widely acknowledged that texts are not just collections of sentences, but have a structure of their own. There has been substantial work in order to account for the different phenomena of *discourse structure* by applying theories of *coherence relations*, e.g. (Mann and Thompson, 1988; Asher, 1993) among others. Coherence relations represent rich semantic linkage (like *Cause*¹ or *Evaluation*) between text segments of varying size.

However, what the minimum size of text segments to relate should be, is still left open to debate. As common approaches argue that coherence relations relate events or situations (e.g. (Hobbs et al., 1993; Asher, 1993)) and that such events are usually introduced

¹Coherence relations in this paper are basically taken from taken Rhetorical Structure Theory (Mann and Thompson, 1988) will appear *emphasized and Capitalized*.

through the means of verbs, it has become standard practice to consider clauses to be the appropriate size for *elementary discourse units*. It has, however, also been observed (Vander Linden and Martin, 1995; Grote et al., 1997) that sometimes phrases may serve as very condensed forms to express elaborate contents. Recently, (Schauer and Hahn, 2000) provided a more detailed analysis when *prepositional phrases* (PPs) may serve as elementary discourse units.

Cursorily viewed, the claims of another recent study stand in contrast to the idea of intra-clausal discourse units: (Schauer, 2000) examined the interplay of coreferential expressions and discourse structure and concluded that referential relations are a good indicator of the discourse structural configurations in case the units examined are entire sentences. This poses the question whether not entire sentences are the appropriate grain size for elementary discourse units.

I will argue that these results – i.e. the different levels of granularity for discourse units – are not incompatible with each other. The approach used in (Schauer and Hahn, 2000) to derive the coherence relation governing a prepositional phrase neatly carries over to the computation of coherence relations signaled by sentence-internal cue-phrases. This then allows an integration with the algorithm using referential relations that was proposed in (Schauer, 2000).

2 Adjuncts as Discourse Units

The question at what size of textual expressions one should start looking for discourse units has not been sufficiently answered yet.

(Mann and Thompson, 1988), for example, state that “Unit size is arbitrary, but the division of the text should be based on some theory-neutral classification. That is, for interesting results, the units should have independent functional integrity. In our analyses, units are basically clauses [...]” [p.248]. It has been noted previously that sometimes phrases already seem to have enough “functional integrity” to give some “interesting results”, however, until recently there has been no detailed proposal what kind of phrases should be considered.

Consider the following (German) example:²

- (1) a. *Mit dem P6LXZ-A versucht Elitegroup neue Kunden zu gewinnen.*
[With the P6LXZ-A — Elitegroup tries to attract new customers.]
- b. *Mit einem PCI-Slot bietet das Motherboard wenig Platz für Erweiterungen.*
[With one PCI slot (only) — the motherboard provides only little space for extensions.]

A straightforward coherence analysis with relations from RST (Mann and Thompson, 1988) takes (1-b) as a single unit and links it to (1-a), probably via an *Evaluation* relation. Paraphrasing sentence (1-b) reveals, however, a plausible decomposition into two interdependent discourse units:

- (2) a. *The motherboard has one PCI slot,*
- b. *so it provides only little space for extensions.*

Clause (2-a) gives an *Explanation* for the subordinate clause (2-b). This might be attributed to the impact of the cue word “so” in (2-b). More generally, it has been claimed that whenever an implicit coherence relation can be made explicit by a paraphrase incorporating a specific cue word, then this coherence relation can always be assumed to hold.³ So, from a methodological point of view it

²In the following, I will summarize the arguments from (Schauer and Hahn, 2000) that focus on prepositional phrases.

³“As a test for the presence of an implicit connection it can be required that the connection could have been explicit...” (Martin, 1992, p.184)

cannot be justified to analyze Sentence (2) as being composed of two elementary units, while the prepositional phrase “with one PCI slot” should be an indistinguishable part of the whole Sentence (1-b).

Obviously, not all kind of phrases may serve as elementary units. As a first criterion to determine whether a phrase is a discourse unit, we propose to consider only those phrases that are not syntactically mandatory *complements* of their governing syntactic head. Complements are assumed to be specified in terms of subcategorization frames, valency lists, etc. *Adjuncts* are phrases which do not match such a complement specification of their governing syntactic head.

In Example (1-b), while it is both possible to omit the optional prepositional phrase “with one PCI slot” and to decompose the sentence such that the prepositional phrase turns into a proper clause, this is only hardly possible for any other phrase in this example, say the subject “the motherboard”.

The second major criterion operates at the semantic level. Semantic specifications of lexemes, independent of the specific semantic theory one subscribes to, are confined to “typical” properties, e.g., events are characterized by agents, patients, instruments, locations, time frames, etc. The interpretation of phrases referring to these attributes can be taken care of by straightforward semantic interpretation mechanisms. As long as the interpretation makes reference only to such “typical” properties it is not necessary to consider such phrases as discourse units. But when an analyst thinks that the interpretation involves (additional) non-typical, unpredictable information, coherence relations may capture their value-adding meaning.⁴ E.g., the prepositional phrase of (1-a) only refers to a such typical property, namely the instrument of an action. This typicality consideration does not carry over to an “explanation” of events, which is the interpretation of “with

⁴The interpretation of complements is assumed to involve only these typical attributes. The interpretation of adjuncts, in contrast, may refer to both typical and non-typical information.

one *PCI slot*" from Sentence (1-b).

Ultimately, however, there is also a third criterion that needs to be considered: namely, that a coherence relation can be found by which the phrase may be related to the remainder of the clause it occurs in (Marcu et al., 1999). We propose that this search for an applicable coherence relation may be performed by trying to paraphrase the phrase and its matrix clause in such a way that it results in two clauses that are related by a cue-phrase.

(Schauer and Hahn, 2000) evaluated the applicability of the criteria by human annotators using two corpora. We found that 11.0% of the prepositional phrases in a German text corpus can be seen as discourse units and 8.0% in an English control set. These PPs are responsible for 14.8% and 12.1%, respectively, of the relations found in the texts. The discourse units triggered by prepositional phrases always occur in satellite (i.e. subordinate) position, which is not surprising, given their status of optional information.

3 From Prepositional Phrases to Coherence Relations

For an automated coherence analysis, we will basically take the preposition as a trigger for computing a coherence relation. To make this discussion more concrete, it is embedded in the framework of SYNDIKATE, a text analysis system under development in our lab (Hahn and Romacker, 1999).

3.1 Semantic Interpretation

Crucial for the discussion here is that the *syntactic analysis* provided by the system allows to distinguish between adjuncts and mandatory complements: for Sentence (1-b), a dependency relation *ppadj* between the prepositional phrase and its governing syntactic head is derived, that results from the valency specification for the main verb "*bietet*" (*provide*). For computing a *semantic interpretation*⁵, the corresponding conceptual correlates of the content words are then checked

⁵We assume the framework of description logics (Woods and Schmolze, 1992).

for role compatibility.

In Example (1-b), the major interpretation constraints derive from the main verb "*bietet*" (*provide*) which is represented by the concept PROVIDE. It has three major conceptual roles, PROVIDE-PATIENT, PROVIDE-CO-PATIENT, and INSTRUMENT. The PROVIDE-PATIENT and PROVIDE-CO-PATIENT roles can be filled by some instance of MOTHERBOARD and SPACE, respectively.

Focusing on the analysis of the PP, each preposition specifies allowed semantic relations (Romacker et al., 1999). In the case of "*mit*" (*with*) they allow an interpretation in terms of the conceptual INSTRUMENT role, so the corresponding role of PROVIDE is filled with *PCI-SLOT*. This causes additional *conceptual interpretation* processes: a production rule checks whether the filler of the INSTRUMENT role may be a PART-OF the conceptual correlate of the syntactic subject, linking SPACE and MOTHERBOARD. Along the line, a HAS-PART-STATE is established that corresponds to the situation that "*the motherboard has one PCI slot*".⁶

3.2 Discourse Interpretation

As discussed above, a preposition does not always signal a coherence relation. Also, a preposition may give rise to different coherence relations depending on context, so one can not simply paste a label of "the" coherence relation into the representation. Instead, *the computation of coherence relations* is performed by testing relation-specific constraints on the configuration of the text-knowledge base resulting from semantic interpretation. Following standard practice, coherence relations are assumed to connect situations. So, we try to derive the relation between the conceptual correlate PROVIDE of the main verb and the inferred state HAS-PART-STATE, which stems from the conceptual interpretation of the prepositional phrase.

⁶In case that the sentence "*the motherboard has one PCI slot*" would have been parsed, the PART-OF relation would have been derived via a production rule operating on the interpretation of the main verb "*has*". Thus, establishing the HAS-PART-STATE in Example (1-b) yields a uniform analysis.

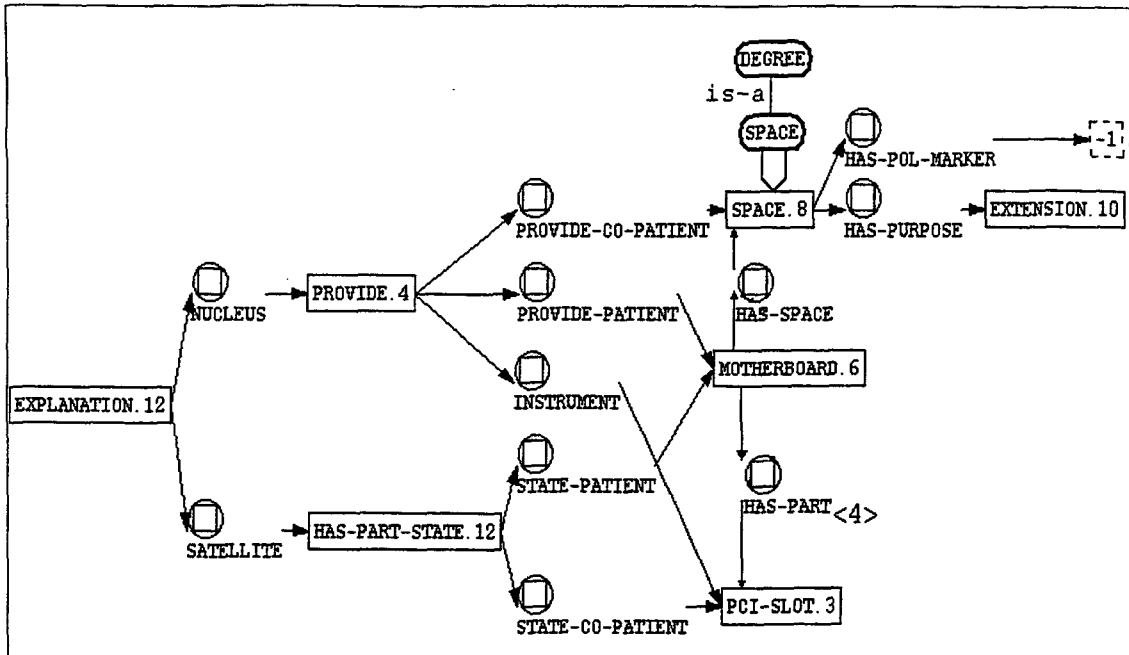


Figure 1: Discourse Interpretation for Example (1-b)

The preposition “mit” (*with*) may signal an *Explanation*. The basis for computing this relation consists in recognizing that the matrix clause gives some kind of evaluation: namely, that the “motherboard provides only little space.” The *degree expression* SPACE, with its associated negative POL-MARKER, (Staab and Hahn, 1997) is the trigger for recognizing the evaluative status of the matrix clause. We can now start looking for the signalled *Explanation* for this negative judgment.

Explanation can be found quite easily: usually, a MOTHERBOARD has up to four PCI-SLOTS (see the default value that is denoted at the PART-OF relation in Figure 1). But actually the situation HAS-PART-STATE is a state in which only one is present, which is obviously “little”. The computations that derive the *Explanation*-relation are all performed by means of production rules.

As a preposition might give rise to different coherence relations, the constraints of all those relations have to be checked. None of the constraints that are checked for other coherence relations signaled by “mit” (e.g. *Con-*

dition) are satisfiable in this example. Similarly, for Example (1-a) no coherence relation would be derived, because the constraints would not be applicable. So, the computational approach is complementary to the semantic criterion of typicality discussed in Section 2.

4 Interclausal Coherence

The proposed mechanism essentially treats the preposition as a kind of cue-phrase that determines a set of possible coherence relations, combining it with an inference mechanism. The main advantage of using such a combination is that it is possible to disambiguate cases in which cue-phrases may give rise to different relations without being confined to testing the complete range of coherence relations. This problem is not restricted to prepositions: (Knott and Sanders, 1998) actually build a multi-taxonomy of cue-phrases in which elements may give rise to several relations, depending on context. E.g., “as” has three connotations: “as” may signal temporal, causal or similarity relationship.

The presented approach carries over neatly to such cue-phrases with multiple relations: indeed, the process described in Section 3 can be seen as a formal reconstruction of the paraphrase. The conceptual representations that is derived for Sentence (2) is basically the same as for (1-b): in both cases a HAS-PART-STATE and a PART-OF relation are computed that connects the MOTHERBOARD and its PCI-SLOT. Finally, the computation of the coherence relation for Sentence (2) differs only with respect to the trigger: in Sentence (2), the cue-phrase “so” triggers the computation instead of the preposition “mit” in (1-b).

The presented approach is thus able to handle coherence relations that operate on the intra-sentential level. Still, this is only a first step to account for the discourse structure of entire texts: cue-phrases are by no means always present. And even if there is a cue-phrase in a new unit to attach, it is still often not clear from the cue alone to which unit the new one should connect:

- (3) a. *Apple presented a new member of the MessagePad organizer family.*
- b. *A new display and a new version of the operating system are introduced.*
- c. *However, the old MessagePad 120 is still available.*

The cue “however” alone does not give enough information to decide whether Sentence (3-c) should connect to (3-b) or to (3-a): further information is needed, like that there is a referential relation between *the old MessagePad 120* and *the MessagePad family*. As a step towards connecting larger units, a recent study (Schauer, 2000) examined the interaction of discourse structure and referential relations among nominal expressions.

4.1 Structural Dependencies

What seems intuitively clear, is that a theory of discourse structure should reflect the structure of the text under consideration. Rhetorical Structure Theory, however, concentrates on the *effects* intended by the author of a text. RST focuses on a “pre-realizational” structure and is not primarily concerned with text

phenomena. The results are analyses that do not account for the *structural constraints in a text*, here referential constraints: the dependency of an anaphoric expression on its antecedent – that it is resolvable and hence interpretable – can hardly be captured in RST. While this is of course not an issue for RST as a theory, it is a prerequisite for any system that wants to account automatically for a text’s discourse structure. As an example consider the following fragment:

- (4) a. *The Vaio F190 comes with convincing equipment.*
- b. *It features a DVD-ROM and a 6 GB harddisk.*

In classical RST, example (4-a) could be said to stand in an *Evaluation* relation to (4-b). The definition of *Evaluation* requires that the satellite evaluates the nucleus, see Figure 2.⁷ However, this would not capture a structural dependency: the pronominal anaphora “it” cannot be interpreted correctly without the antecedent, so (4-b) depends on (4-a).




Figure 2: Structure of *Evaluation*

In RST, one could reflect this referential dependency by analyzing (4-b) as giving *Evidence* for (4-a) (see Figure 3), as an *Evidence* relation may be applied when the hypothesis is the nucleus. Such an analysis would neglect the “convincing” in (4-a) which is quite an explicit evaluation.

In order to account for the structural dependency and the appropriate semantic linkage, we propose to use a new kind of *Evaluation-N*(nucleus) relation. More generally, this loosens the relation between nucleus-satellite assignment and semantic content of coherence relations.

⁷The depicted structures reflect standard RST schemata. The target of the arrow marks the nucleus.




Figure 3: Structure of *Evidence*

4.2 Structural Configurations

Analyzing texts using the outlined notion of structural dependencies as a basic guideline, we determined the structural configurations involving coreference relations. Basically, when two units contain coreferential expressions, they are usually connected by a coherence relation which subordinates the anaphoric unit. In the simplest case, the relation is one of *Elaboration*. However, further linguistic cues or inferences might give rise to semantically “richer” relations.

When more than one coreference relation is involved, issues become more complex. For instance, consider the following example:

- (5) a. *The Panasonic LC90S is a 19"-display.*
 b. *Although its screen size of 482 mm corresponds to a conventional 21"-monitor,*
 c. *considerably less space is required.*
 d. *The device can be attached to a video-card via a USB-connector.*

Obviously, the nominal anaphor “*the device*” in (5-d) requires an antecedent. One possibility would be the pronominal “*it*” in (5-b), leading to a resolution to “*the LC90S*”. However, this is not reflected in the discourse structure that seems most appropriate, cf. Figure 4. The topic of (5-b) and (5-c) (the size) is not further elaborated on in (5-d), so one might say there is a mini-segment boundary between these two sentences. Hence, it would also be not correct to analyze (5-d) as an *Elaboration* of (5a-c), because (5-d) elaborates only (5-a). So, unless further connections (stemming from cue-phrases, inferences or whatever) license a topic continuation, a new unit containing anaphora will connect to the upmost unit bearing their antecedents.

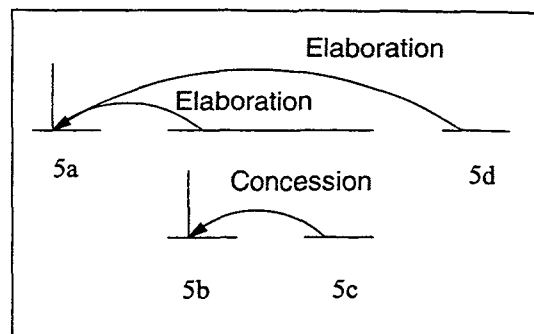


Figure 4: RST analyses for Example (5)

5 From References to Coherence

The configurations just described naturally lead to a combined account of deriving a text’s discourse structure and resolving its referring expressions, repeated from (Schauer, 2000) in Figure 5. Basically, the algorithm uses the successful resolution of anaphoric expressions as a guideline for determining the target unit to which a new unit should be connected which in turn restricts the set of units which should be searched for resolving further referring expressions.

The algorithm assumes the availability of several capabilities for anaphora resolution. First, for a given unit candidate⁸ a set of noun phrases needs to be identified that may be anaphoric expressions. Second, some resolution process is necessary that checks whether an anaphoric expression can be resolved in a given list of possible antecedents. This process makes heavy use of underlying (domain) knowledge (Hahn et al., 1996): in Example (1), “*the P6LXZ-A*” is an instance of a MOTHERBOARD. Since “*the motherboard*” in (1-a) is a definite noun phrase and syntactic as well as conceptual information match with the plausible antecedent “*the P6LXZ-A*”, a referential link can be established, see the IS-COREFERENTIAL relation in Figure 6.

Abstracting away from the actual implementation, the algorithm operates over a tree structure that manages lists of accessible an-

⁸For the sake of argument and because the algorithm is taken from (Schauer, 2000), clauses will be taken as units for the moment. The issue will be discussed below in more detail.

```

tree := tree(centers_forward(first(clauses)), NIL)
clauses := rest(clauses)
forall clause := clauses do
  ana_nodes := array of lists of nodes.
  forall ana_cand := anaphoric_expressions(clause) do
    node := lowest_right_node(tree)
    while node do
      if match(ana_cand, ante_list(node)) then
        ana_nodes[ana_cand] := append(ana_nodes[ana_cand], node)
        node := predecessor(node)
      done
    done
  target_node := find_highest_node_matching_all(ana_nodes)
  /* found at least one antecedent node */
  if target_node then
    /* connect new unit to old node */
    connect(target_node, tree(centers_forward(clause), NIL))
  done

```

Figure 5: Algorithm integrating Coreferences and Discourse Structure

tedecedent candidates and reflects the discourse structure in terms of coherence relations that has been derived so far.⁹

The algorithm now loops through all clauses of a text, building up both the tree and the antecedent lists incrementally. Whenever a new clause has to be considered, its left context consisting of the right frontier of the tree is checked for plausible antecedents.

When all accessible antecedent nodes are determined, the *highest* node that provides antecedents for all resolvable anaphoric expressions in the new unit is taken as the target node, in accordance with the discussion in Section 4.2.

If a new unit contains no referential expression then the algorithm makes no prediction. If the target node has been found, the new unit is connected to it, i.e. the new unit is established as a satellite to the target unit. This means that the new unit opens a new rightmost branch and hence becomes the lowest-right node of the tree. So, the new right frontier consists of the newly attached unit, the

⁹The computed representation of the discourse structure is shown in Figure 6, while the storage of accessible referential objects is managed elsewhere in the system as it interacts with the syntactic parser.

modified node and its predecessors.

The evaluation in (Schauer, 2000) showed that the predictions made by the algorithm depend crucially on the size of the elementary units considered. When clauses are considered to be the elementary unit size, the predictions are correct in up to 81% of the cases in which the algorithm makes a prediction — under the pre-condition that intra-sentential units are handled first. Linking units by referential makes surely no sense when the units considered are phrases: of course, syntactic constraints always mandate that intra-clausal units are related with each other first.

This algorithm is only a useful approximation towards a complete account of a text's discourse structure. Relating to the highest unit providing antecedents works only when there is a mini segment boundary every time an anaphoric expression is used. Although the algorithm ignores other sources of relations between units — which are exactly the cases where the algorithm fails — the results are surprisingly good. As such, enhancing the cue-phrase-only approach by additionally considering coreference relations is in the middle of cheap surface-only approaches and very

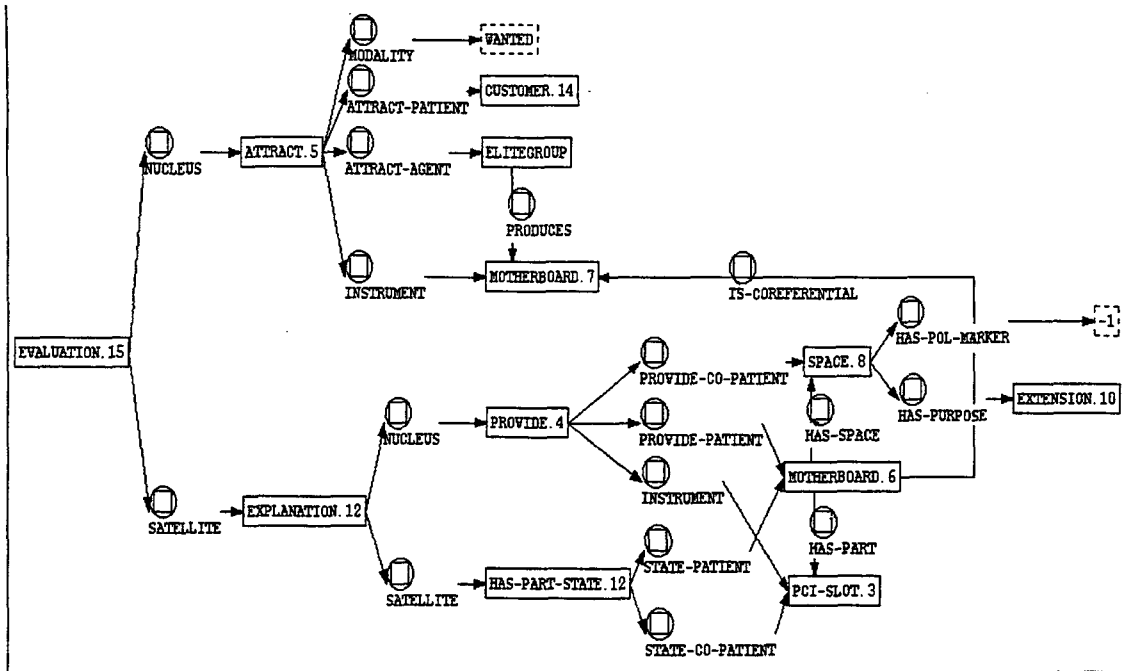


Figure 6: Representation of the Complex Discourse Unit for Example (1)

expensive inference-only methods. A complete approach will also try to account for inter-sentential cue-phrases and also further interconnections, based on the computed representation, see Figure 6.¹⁰

6 Related Work

Rhetorical Structure Theory (Mann and Thompson, 1988) was basically used as the theory on which the presented arguments were grounded. Unfortunately, its definitions are intended to be applied by human annotators and have long resisted approaches to automation because of their vagueness. Previous approaches to (automatically) deriving a text's discourse structure either tried to rely on purely surface-based criteria (cue-phrases, (Marcu, 1998)) or concentrated on inferences about representations of discourse units (Asher, 1993; Hobbs et al., 1993). However, neither of these approaches tries to account for phrases as discourse units. Two re-

cent studies mention the role of PPs as carriers of coherence relations (Vander Linden and Martin, 1995; Grote et al., 1997), but only for the purpose of text generation.

As prepositions (and cue-phrases in general) can signal different coherence relations, the presented computational approach couples a cue-phrase approach like (Marcu, 1998) with inferences using the computed semantic representation. Only recently, there have been some approaches (Webber et al., 1999) that acknowledge the need to incorporate both a treatment of cue-phrase approaches and inferences. However, how these inferences take place is not addressed, neither is the level of intraclausal coherence, that, as we have argued, substantially relies on an integration of both veins. Purely inferential approaches like (Asher, 1993; Hobbs et al., 1993) have also been criticized for being not really applicable because of computational costs. The presented approach tries to minimize these costs by relying on surface-based criteria that reduce the set of coherence relations that have to be considered.

¹⁰The *Evaluation* relation in Figure 6 is due to the evaluative nature of Sentence (1-b), see the discussion in Section 3.

Cue-phrases are not necessarily alone responsible for the discourse structure of texts. Hence, referential relations between nominal expressions were suggested as means for finding attachment points of new units that need to be incorporated. (Corston-Oliver, 1998) enlarges the cue-phrase approach of (Marcu, 1998) in a vein similar to the one presented in Section 5. However, how several coreference relations interact with the resulting discourse structure, is not spelled out. Also, in his approach, it remains unclear how the correct target node to attach to is identified.

Segmented Discourse Representation Theory (SDRT), as described in (Asher, 1993), provides a framework of discourse structure which interacts with referential accessibility constraints. Asher does not rely on coreferences for establishing target units; instead the derivation of a coherence relation (and thereby of the target unit to connect a new unit to) relies on rather abstract connections between “events”. While recognizing coreference relations certainly also requires domain knowledge and inference capabilities, recognizing connections between events seems an even more impossible task. The presented approach is hence more light-weight in nature.

(Webber et al., 1999) apply operations on trees for incorporating cue-phrases in a grammar-driven approach to discourse structure. The presented algorithm could be integrated with their approach, to account for the cases of units that are not connected by referential relations but via cue-phrases.

7 Conclusion

Starting from the question what are the elementary units to consider for a text’s discourse structure, I presented an account for prepositional phrases with adjunct-status as discourse units. Prepositions can be seen as a kind of cue-phrase; however, a preposition does not necessarily signal a coherence relation and even if it does is often ambiguous with regard to the coherence relation signaled. Therefore accounting for prepositional phrases as discourse units requires additional inferences operating on the semantic repre-

sentation of both PP and its matrix clause.

The approach neatly carries over to the phenomena of ambiguous cue-phrases. However, this is still not sufficient to account for deriving the discourse structure of texts in general: cue-phrases are by no means always present and even if there is a cue-phrase, detecting the correct attachment point of a new unit is usually not straight-forward.

As (one step towards) a solution, referential relations between nominal expressions were suggested. The presented algorithm integrates the resolution of anaphora — which also depends heavily on inferences and domain-knowledge — with choosing the target unit to which a new unit should connect to: namely, the highest node providing antecedents to all anaphoric expressions in the new unit. In order for this algorithm to operate successful, it is however necessary that this process is started only after all phenomena of intra-sentential coherence relations have been accounted for, which might be done using the combined approach outlined above.

Returning to the question posed at the beginning of the paper – what is the appropriate size of elementary discourse units – the answer is twofold: first of all, coherence relations can be found to hold between phrases and the clause containing them, so one should indeed start looking for discourse units at the phrase level. However, Syntax requires that the components of sentences group together, and returning to what (Mann and Thompson, 1988) said, sentences have a kind of functional integrity – one that operates on a level that is different from those of phrases. Once this level is reached, larger chunks can be formed, e.g. by referential means.

The presented approach to cue-phrases as well as the use of referential relations will be implemented in the text understanding system SYNDIKATE (Hahn and Romacker, 1999), in order to account for semantically rich relations between larger text chunks and the discourse structure of texts in general. This, however, will require further understanding of the conditions of the coherence relations involved.

Acknowledgments

The author is a member of the Graduate Program on *Human and Machine Intelligence* at Freiburg University, Germany, funded by *DFG*. I would like to thank the staff at the Computational Linguistics Lab, Freiburg, and also the reviewers for their comments.

References

- Nicholas Asher. 1993. *Reference to Abstract Objects in Discourse*. Dordrecht: Kluwer.
- Simon H. Corston-Oliver. 1998. Identifying the linguistic correlates of rhetorical relations. In *Proceedings of the COLING-ACL'98 Workshop 'Discourse Relations and Discourse Markers'*, pages 8–14. Montreal, Canada, August.
- Dan Cristea, Nancy Ide, and Laurent Romary. 1998. Veins theory: A model of global discourse cohesion and coherence. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics*, pages 281–285, Montreal, Canada, August.
- Dan Cristea, Daniel Marcu, Nancy Ide, and Valentin Tablan. 1999. Discourse structure and co-reference: An empirical study. In *The ACL99 Workshop on Standards and Tools for Discourse Tagging*, pages 48–57, June.
- Brigitte Grote, Nils Lenke, and Manfred Stede. 1997. Ma(r)king concessions in English and German. *Discourse Processes*, 24(1):87–118.
- Udo Hahn and Martin Romacker. 1999. SYN-DIKATE – generating text knowledge bases from natural language texts. In *Proceedings of the 1999 IEEE International Conference on Systems, Man, and Cybernetics*, volume 5, pages 918–923. Tokyo, Japan, October.
- Udo Hahn, Katja Markert, and Michael Strube. 1996. A conceptual reasoning approach to textual ellipsis. In *Proceedings of the 12th European Conference on Artificial Intelligence*, pages 572–576. Budapest, Hungary, August.
- Jerry R. Hobbs, Mark E. Stickel, Douglas E. Appelt, and Paul Martin. 1993. Interpretation as abduction. *Artificial Intelligence*, 63:69–142.
- Alistair Knott and Ted Sanders. 1998. The classification of coherence relations and their linguistic markers: an exploration of two languages. *Journal of Pragmatics*, 30(2):135–175.
- William C. Mann and Sandra A. Thompson. 1988. Rhetorical Structure Theory: toward a functional theory of text organization. *Text*, 8(3):243–281.
- Daniel Marcu, Estibaliz Amorrortu, and Magdalena Romera. 1999. Experiments in constructing a corpus of discourse trees. In *Proceedings of the ACL'99 Workshop 'Standards and Tools for Discourse Tagging'*, pages 48–57, University of Maryland, USA, June.
- Daniel Marcu. 1998. A surface-based approach to identifying discourse markers and elementary textual units in unrestricted texts. In *Proceedings of the COLING-ACL'98 Workshop 'Discourse Relations and Discourse Markers'*, pages 1–7. Montreal, Canada, August.
- J.R. Martin. 1992. *English Text. System and Structure*. John Benjamins, Philadelphia/Amsterdam.
- Martin Romacker, Katja Markert, and Udo Hahn. 1999. Lean semantic interpretation. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*, volume 2, pages 868–875. Stockholm, Sweden, August.
- Holger Schauer and Udo Hahn. 2000. Phrases as carriers of coherence relations. In *Proceedings of the 22nd Annual Meeting of the Cognitive Science Society*, pages 429–434, Philadelphia, USA, October.
- Holger Schauer. 2000. Using coreferences for coherence relations. In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics, Student Research Workshop*, Hong Kong, China, October.
- Steffen Staab and Udo Hahn. 1997. Comparatives in context. In *Proceedings of AAAI97 – the 14th National Conference on Artificial Intelligence*, pages 616–621. Providence, USA, July.
- Keith Vander Linden and James H. Martin. 1995. Expressing rhetorical relations in instructional texts: A case study of the purpose relation. *Computational Linguistics*, 21(1).
- Wietske Vonk, Letticia G.M.M. Hustinx, and Wim H.G. Simons. 1992. The use of referential expressions in structuring discourse. *Language and Cognitive Processes*, 3/4(7):301–333.
- Bonnie Webber, Alistair Knott, Matthew Stone, and Aravind Joshi. 1999. Discourse relations: A structural and presuppositional account using lexicalised TAG. In *Proceedings of the 37th Meeting of the Association for Computational Linguistics*, University of Maryland, June.
- William A. Woods and James G. Schmolze. 1992. The KL-ONE family. *Computers & Mathematics with Applications*, 23(2/5):133–177.