DISCOURSE MARKER CHOICE IN SENTENCE PLANNING

Brigitte Grote*
Otto-von-Guericke Universität Magdeburg
Institut für Wissens- und Sprachverarbeitung

Manfred Stede†
Technische Universität Berlin
Projektgruppe KIT

Abstract

In text, discourse markers signal the kind of coherence relation holding between adjacent text spans; for example, *because*, *since*, and *for this reason* are different markers for causal relations. For any but the most simple applications of text generation, marker selection is an important aspect of producing cohesive text. However, present systems use markers in fairly simplistic ways and cannot make use of the full potential of markers that language offers for a given relation. To improve this situation, we propose a specialized lexicon for discourse markers, which holds the relevant constraints and preferences associated with the markers, and which can be used by a text generator to make an informed choice among alternative ways of expressing a relation in the given context. We demonstrate how the lexicon can be employed in the generation process and propose to perform discourse marker choice in the *sentence planning* stage, where the interactions with other generation decisions can be accounted for.

1 Introduction

When a coherence relation ties two adjacent portions of text together, it is often lexically signalled on the linguistic surface with a suitable word—most often a conjunction, but also a preposition, a prepositional phrase or an adverb [Quirk et al. 1972]. The set of words from these grammatically heterogeneous classes that can signal coherence relations we call *discourse markers*. For example, the CONCESSION relation in English can be signalled with the subordinator *although*, the adverb *still*, the preposition *despite*, and a number of other words.

For most coherence relations, language offers quite a variety of such markers, as several studies of individual relations have demonstrated (see references in Section 2). Accordingly, from the generation perspective, a serious choice task arises if the produced text is not only to simply signal the coherence relation, but moreover to reflect pragmatic goals, stylistic considerations, and the different connotations markers have. The importance of these factors was stressed by Scott and de Souza [1990], who gave a number of informal heuristics for when and how to signal the presence of coherence relations in text. Fleshing out the choice task in order to come up with a computational model, though, reveals two sources of difficulty.

For one thing, in addition to syntactic variety, the precise semantic and pragmatic differences between similar markers can be quite difficult to determine. For instance, the CONCESSION markers *although* and *even though* differ merely in emphasis; the CAUSE markers *because* and *since* differ in whether they mark the following information as given or not; the German CAUSE markers *weil* and *denn* differ in the illocution type of the conjuncts (proposition versus statement). Second,

*Complete address: Otto-von-Guericke Universität Magdeburg, IWS/FIN, Universitätsplatz 2, 39106 Magdeburg, Germany; email: grote@iws.cs.uni-magdeburg.de

†Complete address: Technische Universität Berlin, Fachbereich Informatik, Projektgruppe KIT, Sekretariat FR 6-10, Franklinstr. 28/29, 10587 Berlin, Germany; email: stede@cs.tu-berlin.de

the dependencies between marker choice and other generation decisions are rather intricate. The idea of avoiding them is, presumably, the reason for the simplistic treatment of marker choice in typical generators to-date: They regard discourse markers as mere consequences of the structural decisions, hence do not perform any choice. We wish to demonstrate, however, that this strategy, which is typical for dealing with closed-class lexical items in general, is too great a simplification in these cases.

In this paper, we propose to use a discourse marker *lexicon* as a declarative resource for the *sentence planning* stage of the text generation process. The paper is organized as follows. Section 2 examines the role of discourse markers in NLG and reviews the state of the art. Section 3 briefly summarizes the ideas on sentence planning that have arisen in the past few years and argues that for a sophisticated treatment of discourse marker choice, a dedicated lexicon is to be used as one information resource in sentence planning. Section 4 introduces the discourse marker lexicon we are currently developing, and Section 5 describes how this lexicon can be usefully employed in the sentence planning phase to realize more flexible marker production.

2 Discourse markers in NLG

We follow Moser and Moore [1995] in assuming that three distinct though interrelated decisions have to be made when generating discourse markers: Whether to place a marker or not (*marker occurrence*), where to place a marker (*marker placement*), and finally, which marker to use (*marker selection*).

Research on connectives in the context of NLG has focused on the *selection* of markers to produce coherent and cohesive multi-sentential text. Studies fall into two distinct groups: First, studies are concerned with identifying the characteristic properties of a small set of similar markers, and determining the reasons behind choosing a particular marker from this set in a given context; examples are the markers *since* and *because* [Elhadad and McKeown 1990], or the temporal markers *before* and *while* [Dorr and Gaasterland 1995]. Second, a number of studies take particular (RST-)relations as a starting-point, and examine how these relations are signalled on the linguistic surface; examples are the PURPOSE, RESULT and PRECONDITION relations [Vander Linden 1994], the CONCESSION relation [Grote et al. 1997], and the subject-matter relations occurring in a technical domain [Rösner and Stede 1992, Delin et al. 1996]. However, these are all isolated studies, geared towards a particular application. There is at present no overall framework that supports informed and motivated marker generation for more than a small set of markers and relations. The broadest overview on discourse markers to our knowledge is the descriptive work of Knott and Mellish [1996], but it does not specifically address the NLG perspective. Moser and Moore [1995] and DiEugenio et al. [1997] also take a broader view on marker production in that they try to determine general factors that influence the use of markers in text, and in that they consider more than pairs of propositions. However, they are largely concerned with marker occurrence and placement, not with marker selection.

3 Sentence planning

The traditional split of NLG systems in a content determination/what-to-say component and a realization/how-to-say component was in recent years supplemented by an intermediate stage: sentence planning, sometimes called *micro-planning* (e.g., Rambow and Korelsky [1992]). The primary motivation for this step is to relieve the text planner from language-specific knowledge, and to relieve the realization component from any planning or decision-making that potentially

affects the meaning of the utterance. Hence, better control of the overall generation process is gained. We do not elaborate the advantages further here; see, for example, [Panaget 1994].

What are the specific decisions to be made by the sentence planner? We think it is important to separate the *formative decisions* from the *motivations* that lead to the particular choices. Following Wanner and Hovy [1996], a sentence planner has to make the following decisions: Fine-grained discourse structuring, including discourse marker choice; sentence grouping and sentence content determination; clause-internal structuring; choice of referring expressions; lexical choice.[1] Two groups of considerations are important for these tasks: First, the motivating factors such as stylistic choices, semantic relations, intentions, theme development, focusing, discourse history. Second, the interactions with other desicions, because different formative decision may realize the same motivation. In contrast to present NLG systems, which realize the production of marker choices as a mere consequence of other sentence level decisions, we think that sentence planning interactions ought to be respected for discourse markers, too, as the following examples illustrate:

- Ordering of related clauses (cause-effect vs. effect-cause)
  Because *he was unhappy, he asked to be transferred.* vs. *He asked to be transferred,* **for** *he was unhappy.* vs. * **For** *he was unhappy, he asked to be transferred.*
- Aggregation
  He has quarrelled with the chairman. He resigned from his post. vs. *He has quarrelled with the chairman* **and** *resigned from his post.*
- Delimit individual sentences
  They fought a battle. **Afterwards,** *it snowed.* vs. **After** *they fought a battle, it snowed.*
- Clause-internal structuring (hypotaxis vs. parataxis)
  Although *he has tried hard, he failed.* vs. *He tried hard,* **but** *he failed.*
- Lexical choice (*to know* vs. *ignorance*)
  She died, **because** *she didn't know the rules.* vs. *She died* **through** *ignorance of the rules.*
- Realizing negation
  He will not attend **unless** *he finishes his paper.* vs. *He will attend* **if** *he finishes his paper.*

Due to these interdependencies, any fixed order of making decisions in sentence planning will impose limitations on the expressiveness of the system. Accordingly, we advocate a flexible order of decision-making, as it can be realized in a blackboard-based architecture such as proposed by DIOGENES [Nirenburg et al. 1989] and HealthDoc [Wanner and Hovy 1996]. Moreover, the individual modules or knowledge sources should rely on declarative representations as much as possible; otherwise the control process becomes extremely complicated. And one of the declarative sources of information, we feel, should be a lexicon that assembles specifically the information associated with discourse markers.

4 The Discourse marker lexicon

4.1 Discourse markers as lexical entities

The traditional distinction between content words and function words (or open-class and closed-class items) relies on the stipulation that the former have their "own" meaning independent of the

[1] How exactly these tasks are to be accomplished depends on the nature of the input and output representations, and thus on the architecture of the generator. In Section 5, we will introduce the framework we are using and characterize the integration of marker choice into the sentence planning phase.

context in which they are used, whereas the latter assume meaning only in context. Then, content words are assigned to the realm of the lexicon, whereas function words are treated as a part of grammar. For dealing with discourse markers, we do not regard this distinction as particularly helpful, though. These words can carry a wide variety of semantic and pragmatic overtones, which render the task of selecting a marker meaning-driven, as opposed to a mere consequence of structural decisions.

Furthermore, notice that a number of lexical relations customarily used to assign structure to the universe of "open class" lexical items can be applied to discourse markers as well. For instance, the German words *obzwar* and *obschon* (both more formal variants of *obwohl* = *although*) are at least very close to being **synonyms**. As for **plesionyms** (near-synonyms), *although* and *though*, according to Martin [1992], differ in formality, and *although* and *even though* differ in terms of emphasis. *If* and *unless* can be seen as **antonyms**, as they both express conditionality, but with opposite polarity. Some markers are more specific than others, thus display **hyponymy**. E.g., *but* can signal a general CONTRAST or a more specific CONCESSION. Finally, other than being more or less specific, some markers can signal quite different relations; e.g., *while* can be used for TEMPORAL CO-OCCURRENCE, and also for CONTRAST. Hence, the marker is **polysemous**.

For these reasons, discourse markers should be described by a dedicated lexicon that provides a classification of their syntactic, semantic and pragmatic features and characterizes the relationships between similar markers. This will be a lexicon whose main grouping criterion is *function* rather than grammatical category; not surprisingly, this is motivated by the production perspective, where the parameters governing the generation decisions play the central role.

4.2 Methodology

Methodological considerations pertain to the two tasks of determining the set of words we regard as discourse markers, and determining the lexical entries for these words.

Finding the "right" set of discourse markers is not an easy task, since the common lexicographic practice of having syntactic behaviour as the criterion for inclusion does not apply. Knott and Mellish [1996] provide an apt summary of the situation. Their 'test for relational phrases' is a good start, but geared towards the English language (we are investigating German as well), and furthermore it catches only items relating clauses; in *Despite the heavy rain, we went for a walk* it would not detect a cue phrase. To identify more markers, we worked with traditional dictionaries and with grammars like Quirk et al. [1972] and Helbig and Buscha [1991]. The resulting set of markers is further validated by investigating coherence relations and their possible realizations; here, we can draw on our earlier work [Rösner and Stede 1992, Grote et al. 1997].

As for the shape of the lexical entries, there are two tasks: First, determining the distinguishing features and classifying markers according to these features, and second, finding appropriate computational representations. At present, we are mostly concerned with the first step, but in section 5, we make an initial proposal for representations.

Regarding the set of features, our goal can be characterized as finding a synthesis of two different perspectives on marker description, between which there has been little overlap in the research literature: *Text linguistics* considers markers as a means to signal coherence, and provides us with insights on the semantic and pragmatic properties of marker *classes*, hence approaches the matter "top-down". On the other hand, *grammars* and *style guides* provide syntactic, semantic and stylistic properties of individual markers, thus look "bottom-up".

Specifying the distinctions within sets of similar markers can be quite subtle. In addition to drawing on our earlier work cited above, we employ techniques such as paraphrasing, Knott's substitution test [Knott and Mellish 1996], analysis of typical distributions using corpora, and contrastive studies. Extracting features in this way seems justified since at this stage we are—unlike

feature	*unless*	*for*	*however*	*even though*	*notwithstanding*
syntactic					
part-of-speech	subordinator	subordinator	adverb	subordinator	preposition
connection type	hypotaxis	hypotaxis	intersent.	hypotaxis	intraclausal
scope of marker	S simple	any	any	S simple	N and S simple
position	front	front	front/mid/end	front	front
linear ordering	N S	N S	S N	any	any
semantic					
semantic relation	condition	cause	concession	concession	concession
polarity	act negated	any	any	any	any
functional order	any	effect-cause	conceding-conceded	any	any
nuclearity	N: act	N: effect	N: conceded	N: conceded	N: conceded
pragmatic					
formality	standard	standard	standard	standard	formal
emphasis	none	none	none	intensified	none
discourse relation	PRECONDITION	VOL.CAUSE	CONCESSION	CONCESSION	CONCESSION

Table 1: *Sample lexicon entries for some English markers*

DiEugenio et al. [1997]—not concerned with the predictive power of individual features but rather with decomposing markers into features that are relevant for integrating marker choice into sentence planning.

4.3 The shape of the lexicon

The initial set of features we have thus obtained can be grouped in the traditional categories:

Syntactic features are the *part-of-speech* of a marker and the type of *connection* it establishes (prepositions link constituents within a clause; conjunctions build a paratactic or hypotactic structure, but some can also function as intersentential linkers). The *scope* of a marker is the complexity of the segments it can combine (complex subtree or simple propositions). The *linear ordering* of the conjuncts can differ from marker to marker (e.g., with the connective *for*, the subordinate clause is always postponed) as well as the marker's *position* within the segment (e.g., prepositions always occur at the beginning of a segment; adverbs like *however* can occur in front-, mid- and end-position).

Semantic features are foremost the *semantic relation* established (e.g. causal or temporal link). Some markers show a particular behaviour towards *negation*, which is related to *polarity* (e.g., *if* versus *unless*). Further, we observe that certain markers impose what we term a *functional ordering*, for instance, *for* requires the order *effect-cause*.

Pragmatic features include the *discourse relation* expressed by the marker and the type of *illocutionary acts* it conjoins (e.g., German *weil* links propositions, *denn* links judgements). Some markers differ in terms of *presuppositions* and the assignment of *given/new* (e.g., *because* versus *since*). Stylistic features represent dimensions like *formality* and *emphasis*.

To illustrate how these features discriminate between markers, table 1 gives five preliminary sample entries. N is a shorthand for nucleus in the RST sense, S for satellite. Notice that table 1 is merely an illustration, and not all the features introduced above are actually required for classification. Issues of lexical representation, including dealing with polysemy and homonymy in some inheritance-based formalism, will be addressed in a later stage of the project.

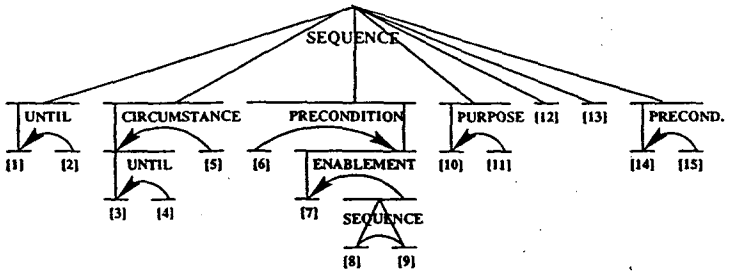

Figure 1: *Input structure to sentence planner*

5 The discourse marker lexicon in sentence planning

Having outlined the discourse marker lexicon as a general resource, we now turn to the question of using it in sentence planning. Even though the lexicon is still under development, we will illustrate with several prototypical representations how a sentence planner can exploit the various realization options offered by the lexicon.

We assume the following framework: a discourse structure tree loosely based on RST [Mann and Thompson 1988] serves as input to the sentence planner. RS-trees comprise a set of propositions as leaf nodes; the internal nodes represent coherence relations holding between the daughter nodes. The tree is encoded in the description logic LOOM [MacGregor and Bates 1987], and the propositions are represented following the ontology used in the MOOSE system [Stede 1996]. The nature of these representations need not concern us here, but it is important that they are all "grounded" in a knowledge base so that type checking via subsumption can take place.

The output of the sentence planning module is a sequence of lexicalised sentence-semantic specifications (SemSpecs), based on SPL [Kasper 1989]. Accordingly, sentence planning in this framework amounts to *linearizing a discourse representation tree.* As front-end sentence generator, we use KPML [Bateman 1997]. A sample input structure from the domain of maintenance manuals is given in figure 1; figure 2 shows one possible realization. Numbers in the tree correspond to text segments, and each segment corresponds to one underlying proposition.

> [Wait]1 until [the engine is cool]2, then [turn the radiator cap clockwise]3 until [it stops]4. [DO NOT PRESS DOWN WHILE TURNING THE CAP]5. After [any remaining pressure has been relieved]6, [remove the cap]7 by [pressing down]8 and [again turning it counterclockwise]9. [Add enough coolant]10 to [fill the radiator]11, and [reinstall the cap]12. [Be sure to tighten it securely]13. [Fill the reserve tank up to the max mark]14 with [the engine cold]15.

Figure 2: *One linguistic realization of the RST-tree*

5.1 The "generation view" of the discourse marker lexicon

From the production perspective, the lexical features are to be classified with respect to when and where they come into play in the generation process; this amounts to one particular "view" on the information coded in the lexicon. We propose these categories:

- **Applicability conditions**: The necessary conditions that need to be present in the input representation for the marker to be a candidate. Chiefly, this is the semantic/discourse relation to be expressed, and also (if applicable) features pertaining to presuppositions and intentions.

- **Combinability conditions**: The constraints that the marker imposes on its neighbouring linguistic constituents (the 'syntagmatic' dimension). These are syntactic constraints on subcategorization and semantic type constraints, which interact with other realization decisions in sentence planning.
- **Distinguishing features**: If preferential choice dimensions, such as style, brevity, etc., are attended to in the system, then these features serve to distinguish markers that are otherwise (nearly) synonymous (the 'paradigmatic' dimension).

For encoding this information, we adopt the framework used in the lexicalization approach of the MOOSE sentence generator [Stede 1996]. Here, lexicon entries consist of (*inter alia*) the three zones *denotation*, *partial SemSpec* (PSemSpec), and *stylistic features*. The denotation is the part to be matched (qua subsumption) against the input representation; it may contain type restrictions. The PSemSpec is an SPL-like template that includes a `:lex` annotation with the actual lexeme and possibly variables that are replaced by other PSemSpecs in the course of the lexicalization process. Also, any realization directives needed by the front-end generator are stated here. Stylistic features are used for preferential choice between words that would all be applicable in a particular context.

When generalizing this framework to include discourse markers (and hence allowing for producing complex sentences), the denotation of a marker would be an RST relation[2] with variables for the relata, possibly enriched with type constraints. For relations with a nucleus and a satellite, we always write them in this order, hence `(RELATION NUCLEUS SATELLITE)`. As a simple case, consider the subordinating conjunction *until*, which we take to be a marker of the relation UNTIL[3], a straightforward case indeed. Its denotation is `(UNTIL X (STATE Y))`, meaning that it can be used to verbalize any `UNTIL` node whose satellite is of type `STATE`, according to the ontology or domain model in the knowledge base.

The variables used in the denotation also appear in the PSemSpec of *until*, so that partial SemSpecs can be combined together correctly. Here, the nucleus of the UNTIL relation becomes the `domain` of the `rst-until` relation as defined in the KPML Upper Model,[4] and the satellite is mapped to `range`, which we further constrain to be a `relational-process` (in Upper Model terms). Furthermore, we add `:theme X` to ensure that the nucleus is ordered before the satellite (to avoid *until Y, X*). The complete lexicon entry together with a few more examples is given in figure 3: The denotations and PSemSpecs for the subordinating conjunctions *until* marking UNTIL and *after, if, then, unless* marking PRECONDITION, and for the preposition *with* in its function as marker for the relation PRECONDITION.

5.2 Marker choice

In MOOSE, lexical options constitute the search space for building a lexicalized semantic sentence specification. Now, we generalize this idea to discourse trees: For propositional nodes MOOSE calculates all possible lexical options; for coherence relation nodes, the list of options realizing the node is taken from the discourse marker lexicon by matching the node against the applicability conditions of the lexicon entries. Thus, the entire discourse tree is annotated with verbalization options, which together constitute the search space for sentence planning.

[2]The relations used in denotations effectively constitute the interface between the lexicon and the text planner producing the discourse tree. At present we use RST, but we regard this only as an interim solution. For the purposes of this paper, the precise inventory of relations used is not critical.

[3]The discourse relation UNTIL was introduced by Rösner and Stede [1992]; its status is somewhat questionable, but since we do not address the general issue of discourse relations here, we simply assume its existence.

[4]All Upper Model discourse relations bear the prefix 'RST', which at this point unfortunately might produce confusion about the variety of relations under discussion.

```
UNLESS
:DEN  (PRECONDITION X Y)
:PSEM (r / rst-precondition
         :domain (p / (process X)
                    :polarity negative)
         :range Y
         :theme X)

THEN
:DEN  (PRECONDITION X Y)
:PSEM (X / process)
      (Y / process
         :conjunctive precondition
         :lex then)

IF
:DEN  (PRECONDITION X Y)
:PSEM (r / rst-precondition
         :domain X
         :range (p / (process Y)
                   :tense present)
         :lex if)
```

```
UNTIL
:DEN (UNTIL X (state Y))
:PSEM (r / rst-until
         :domain X
         :range (Y / relational-process)
         :theme X
         :lex until)

WITH
:DEN  (PRECONDITION X  (state Y))
:PSEM (p / (process X)
         :precondition (r / (property-ascription Y)
                          :lex with))

AFTER
:DEN  (PRECONDITION X (activity Y))
:PSEM (p / posterior
         :domain X
         :range (p / (process Y)
                   :tense [perfect, past perfect])
         :lex after)
```

Figure 3: *Sample (partial) lexicon entries*

To illustrate this approach, consider the propositions 14 and 15 in the sample text: *Fill the reserve tank with the engine cold.* Here, the PRECONDITION relation is signalled by the intraclausal linker *with* (see the lexicon entry above). Other realizations of this RS subtree are ([Vander Linden 1994] offers a similar range):

1. **If** *the engine is cold, fill the reserve tank up to the max mark.*
2. **When** *the engine is cold, fill the reserve tank up to the max mark.*
3. *Fill the reserve tank up to the max mark,* **only if** *the engine is cold.*
4. **After** *the engine has cooled down, fill the reserve tank to the upper mark.*
5. *Do not fill the reserve tank (up to the max mark)* **unless** *the engine is cold.*
6. *Make sure that the engine is cold.* **Then**, *fill the reserve tank up to the max mark.*

To arrive at variant formulations of this kind, depending on different parameters and/or context, our first step is to set up the search space of verbalization options. While MOOSE performs this step for propositions, we will here focus on the coherence relation nodes. In our example, the marker lexicon yields a set of markers that match the applicability condition (PRECONDITION X Y): *after, if, only if, then, unless, when* and *with*. These are annotated at the node, as shown in figure 4, where the leaf nodes are annotated with (shorthands for) some of the lexical options found by MOOSE.

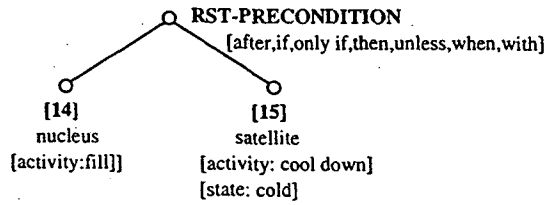

Figure 4: *Annotated subtree for [14] [15]*

From this search space, different decisions made by sentence planning "expert" modules lead to different verbalizations. For instance, assume that the sentence-structuring expert calls for a hypotactic structure; this is satisfied by PSemSpecs of the form: `(r / rst-precondition :domain X :range Y)`, hence by the markers *if, only if, unless* and *when*. If the clause-ordering expert calls for the order satellite–nucleus, *unless* is ruled out as it requires the nucleus to be stated first (see the lexicon entry below). The remaining choice between *only if, if* and *when* is left to fine-grained discrimination (e.g., *only if* is more emphatic), which we do not elaborate here.

Alternatively, assume that the sentence-delimitation expert posits that the relation be expressed in two separate sentences. As a consequence, the ordering is satellite–nucleus. These constraints are satisfied only by the marker *then* (example 6). The sequence of PSemspecs associated with the marker then further constrains the other sentence planning decisions (see figure 3).

Now, it might also be the case that the lexicalization expert (e.g., MOOSE) calls for verbalizing the *result* of the cooling process (proposition [15]) only and proposes the lexeme *be cool*. Now, the marker *after* is out, as it requires the satellite to be realized as a subordinate clause with a process of type `activity` (see lexicon entry of AFTER in figure 3). Alternatively, if the lexeme chosen is *cool down*, markers such as *with* are not available, as its PSemSpec allows for combining with a `property-ascription` only. Now, if some other expert decides to use a negation with the nucleus, *unless* is selected as marker since it expects a negative polarity in the nucleus; its (partial) lexicon entry is shown figure 3.

Selecting *unless* in turn restricts the options for other sentence planning tasks, since its PSemSpec states that a hypotactic structure with the subordinate clause in sentence-final position is needed (due to the `:theme X` line). In short, decisions can be propagated in both directions: from other formative decisions to marker choice, and from marker choice to other decisions. Imagine that the process of tree linearization be driven by the overall goal of producing concise text; in this case, the flexibility in ordering decisions allows for producing short text by choosing *with* and letting the other decisions follow.

We have characterized a constraint-based mechanism that does not impose a strict order on making decisions in linearizing the discourse tree. Various ways of implementing such a scheme can be imagined; one is the blackboard-based approach suggested by Wanner and Hovy [1996], another is the "Hunter-Gatherer" search paradigm introduced by Beale [1996].

6 Summary

Present-day text generation systems typically employ quite simplified approaches for signalling discourse relations in text. Our work aims at enabling generators to truly *choose* discourse markers on the basis of generation parameters and context. This way, we gain variety in marker usage that is not just random but controlled. Furthermore, we are interested in a uniform, declarative representation of the information necessary. Approaches that encode marker choice in the grammar (such as Vander Linden [1994]), while certainly an improvement over previous 1:1 mappings between relations and markers, loose flexibility when it comes to account for the interactions between marker choice and other sentence planning decisions.

These considerations led us to develop a lexicon of discourse markers. While its construction is still in progress, we have shown samples of the kind of representations we are aiming at, and we have demonstrated how such lexicon entries can be employed as a resource in the sentence planning phase of the generation process. In our framework, lexicon entries consist of applicability conditions (for deciding whether the marker is a candidate at all), partial sentence semantic specifications (for combining the marker with neighbouring constituents), and distinguishing features for paradigmatic choice. We have described how an input discourse structure tree can be linearized into a sequence

of sentence plans, given a sentence planner that exploits the information supplied by a discourse marker lexicon.

References

[Bateman 1997] Bateman, J. 1997. Enabling Technology for Multilingual Natural Language Generation: The KPML Development Environment. In *Journal of Natural Language Engineering*, 3(1), 15-55.

[Beale 1996] Beale, S. 1996. Hunter-Gatherer: Applying Constraint Satisfaction, Branch-and-Bound and Solution Synthesis to Computational Semantics. NMSU Technical Report, MCCS-96-289.

[Delin et al. 1996] Delin, J., D. Scott, A. Hartley. 1996. Pragmatic Congruence through Language-Specific Mappings from Semantics to Syntax. In *Proc. of the 16th Conference on Computational Linguistics*, Copenhagen, 292-297.

[DiEugenio et al. 1997] DiEugenio, B., J. Moore, M. Paolucci. 1997. Learning Features that Predict Cue Usage. In *Proc. of the 35th Conference of the Association for Computational Linguistics*, Madrid, 80-87.

[Dorr and Gaasterland 1995] Dorr, B., T. Gaasterland. 1995. Selecting Tense, Aspect and Connecting Words in Language Generation. In *Proc. of the 14th International Joint Conference on Artificial Intelligence*, Montreal, 1299-1305.

[Elhadad and McKeown 1990] Elhadad, M., K.R. McKeown. 1990. Generating Connectives. In *Proc. of the 13th Conference on Computational Linguistics*, Helsinki, 97-101.

[Grote et al. 1997] Grote, B., N. Lenke, M. Stede. 1997. Ma(r)king Concessions in English and German. In *Discourse Processes* 24(1), 87-118.

[Halliday and Hasan 1976] Halliday, M.A.K., R. Hasan. 1976. *Cohesion in English*. London: Longman.

[Helbig and Buscha 1991] Helbig, G., J. Buscha. 1990. *Deutsche Grammatik: Ein Handbuch für den Ausländerunterricht*. Berlin, Leipzig: Langenscheidt, Verlag Enzyklopädie.

[Kasper 1989] Kasper, R. 1989. A Flexible iInterface for Linking Applications to Penman's Sentence Generator. In *Proc. of the DARPA Workshop on Speech and Natural Language Processing*, University of Pennsylvania, 153-158.

[Knott and Mellish 1996] Knott, A., C. Mellish. 1996. A Feature-based Account of the Relations Signalled by Sentence and Clause Connectives. In *Language and Speech* 39 (2-3).

[MacGregor and Bates 1987] MacGregor, R., R. Bates. 1987. The Loom Knowledge Representation Language. Technical Report ISI/RS-87-188, USC/ISI.

[Mann and Thompson 1988] Mann, W., S. Thompson. 1988. Rhetorical Structure Theory: Towards a functional theory of text organization. In *TEXT*, 8(2), 243-281.

[Martin 1992] Martin, J. 1992. *English Text – System and Structure*. Amsterdam: John Benjamins.

[Moser and Moore 1995] Moser, M., J. Moore. 1995. Using Discourse Analysis and Automatic Text Generation to Study the Discourse Cue Usage. In *Proc. of the AAAI Spring Symposium on Empirical Methods in Discourse Interpretation and Organization.*

[Nirenburg et al. 1989] Nirenburg, S., V. Lesser, E. Nyberg. 1989. Controlling a Language Generation Planner. In *Proc. of the 11th International Joint Conference on Artificial Intelligence*, Detroit.

[Panaget 1994] Panaget, F. 1994. Using a Texual Representation Level Component in the Context of Discourse and Dialogue Generation. *Proc. of the 7th International Workshop on Natural Language Generation*, Kennebunkport, Maine, 127-136.

[Quirk et al. 1972] Quirk, R., S. Greenbaum, G. Leech, J. Svartvik. 1992. *A Grammar of Contemporary English*. Harlow: Longman, (20th ed).

[Rambow and Korelsky 1992] Rambow, O., T. Korelsky. 1992. Applied Text Generation. In *Proc. of the Conference on Applied Natural Language Processing*, Trento.

[Rösner and Stede 1992] Rösner, D., M. Stede. 1992. Customizing RST for the Automatic Production of Technical Manuals. In R. Dale et al. (eds.) *Aspects of Automated Natural Language Generation*. Berlin: Springer, 199-214.

[Scott and de Souza 1990] Scott, D., C. de Souza. 1990. Getting the Message across in RST-based Text Generation. In R. Dale et al. (eds.) *Current Research in Natural Language Generation*. Academic Press, 31-56.

[Stede 1996] Stede, M. 1996. Lexical Paraphrases in Multilingual Sentence Generation. In *Machine Translation* 11, 75-107.

[Vander Linden 1994] Vander Linden, K. 1994. Generating Precondition Expressions in Instructional Text. In *Proc. of the 15th Conference on Computational Linguistics*, Kyoto.

[Wanner and Hovy 1996] Wanner, L., E. Hovy. 1996. The HealthDoc Sentence Planner. In *Proc. of the 8th International Workshop on Natural Language Generation*, Herstmonceux Castle, 1-10.