

# INTRODUCING MAXIMAL VARIATION IN TEXT PLANNING FOR SMALL DOMAINS

Erwin Marsi \*

Department of Language and Speech  
University of Nijmegen

## Abstract

This work describes a method for text planning that is suitable to small domains like train table information. Our aim is to introduce maximal variation in the packaging of information and in the linear order of its presentation. To this end, we regard text planning as a goal-driven process that dynamically constructs a text plan. The goal is a state where all information in the input is shared with the user; the means to achieve this goal are utterances. The application of utterances is limited by constraints that refer to the user's current state of knowledge. This approach to text planning can be conveniently implemented as a Functional Unification Grammar. In addition, we show how optional or inferable information can be accounted for, how focus can be distributed, and how the generation of anaphoric expressions can be constrained by looking at the form and content of a previous utterance.

## 1 Introduction

This work on text planning is part of a project that is concerned with investigating Dutch prosody by implementing a concept-to-speech system. The project focuses on the prosodic module, which predicts the pitch accents and the prosodic boundaries of an utterance on the basis of its semantic and syntactic structure and its discourse context. The key idea is that a natural language generator, as opposed to a parser, generates extensive and reliable information about the linguistic structure of an utterance, and is therefore particularly suitable to provide input to the prosodic module. This approach requires at least two things from the generator. First, it should generate all information that the prosodic module needs for deriving the prosodic structure of an utterance. Second, it should generate as much variation as possible, in order to put the prosodic module to the test. Given a conventional architecture consisting of a text planner followed by a surface generator, these requirements affect the text planner. For instance, it should keep track of the information status of concepts, because the distinction between old and new information is important for pitch accent placement. With respect to the second requirement, it should be able to paraphrase one and the same conceptual structure as different semantic structures, which are in turn realized as different sentences by the surface generator.

This paper describes a text planner that meets these requirements. It is described on the basis of an application of concept-to-speech in which train table information is taken as input to generate a spoken description, in Dutch, of how to get from one place to another by train. The approach, however, is easily adaptable to similar domains. Since we are primarily interested in generating linguistically rich and maximally varied input for the prosodic module, the text planner is rather uncomplicated and ignores many other aspects of text planning like rhetorical

---

\*Thanks to Peter-Arno Coppen, Wim Claassen, Carlos Gussenhoven, and two anonymous reviewers for their useful comments and corrections.

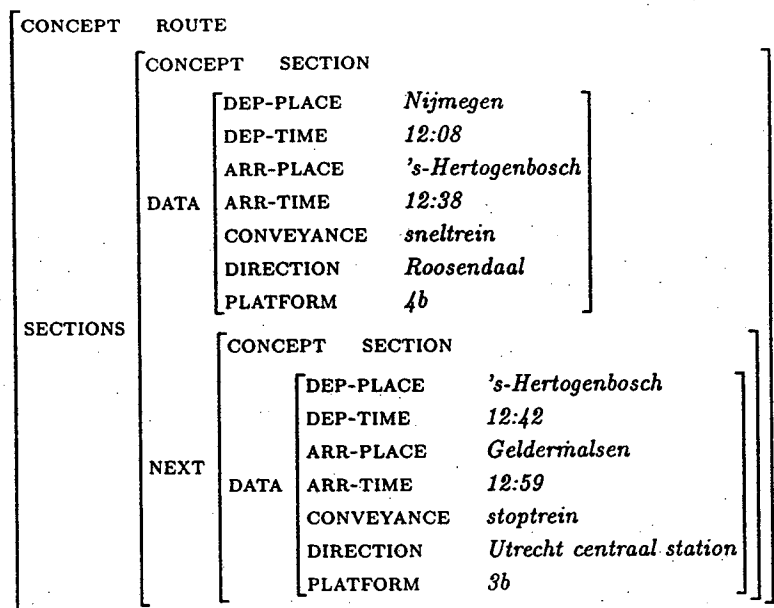


Figure 1: Example of an input structure

structuring of the text or tailoring information to the user. In fact, there is no real dialogue with the user in the sense that the system is capable of reacting on feedback from the user. Also, efficiency considerations (real time behaviour) have not played a role. The interesting points, however, are that the text planner employs a constraint-based approach to produce variation and that its implementation is completely grammar-based within the framework of Functional Unification Grammar.

## 2 A Functional Unification Grammar for text planning

### 2.1 Input structures

The input to the text planner comes from an existing train travel information system. In response to a query typed by the user, it outputs travel information in a tabular format. This information is mapped to a feature description (FD) of hierarchically structured concepts in a straightforward way. For instance, the FD in Figure 1 represents a journey with one change. The top concept, representing the whole journey from departure place to arrival place, is called ROUTE.<sup>1</sup> It is composed of one or more SECTION nodes, each of which represents a partial journey from one place to another. A section node is accompanied by information about the place and time of departure, the place and time of arrival, the type of conveyance, its direction, and the platform it leaves from. Notice that the attribute NEXT serves as a link to the subsequent section.

### 2.2 Text planning grammar

Text planning is regarded as the process of mapping the input structure to a sequence of semantic structures, which will ultimately be realized as spoken utterances. Evidently, not

<sup>1</sup>It also contains information like the total amount of traveling time and the number of changes, which is used to generate a summary of the journey. This option will be ignored here.

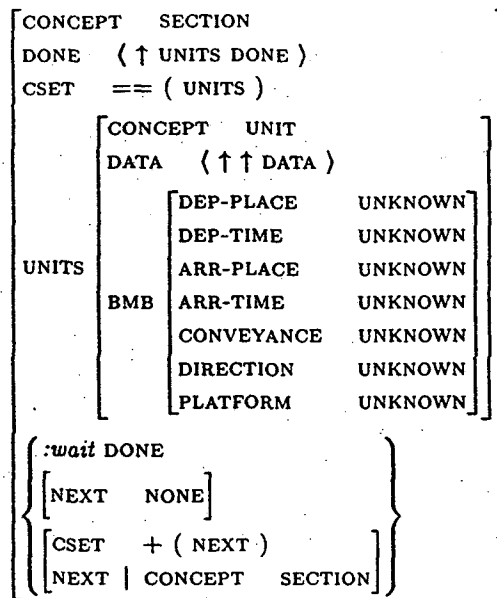


Figure 2: Grammar alternative for a section

all the information in the input can be expressed in a single utterance, so the text planner must divide it into smaller packages. The information within a package should be coherent and the linear order of the packages should make sense. For instance, it is quite odd to start the description of a section with the arrival place and arrival time, that is, without mentioning the departure place and departure time first. Ruling out certain ways of information packaging is of course a matter of common sense; it is always possible to come up with a context in which a very marked order of presentation is acceptable. The obvious solution is to use one or more templates that prescribe acceptable ways of presenting the information. However, as explained above, our goal is to generate as much variation as possible. Using just a limited number of templates would severely restrict the amount of variation at the level of text planning. To obtain more variation, one has to create an extensive list of templates, which accounts for all possible ways of packaging and linear ordering of information.

The alternative is to adopt a dynamic approach to text planning, and to consider it as an attempt to achieve a particular goal under certain constraints (Hovy 1991). The goal is a transfer of all available information, i.e. a state where the user knows all the information that is in the input structure. This does not imply, however, that all data have to be explicitly expressed, because the listener may infer some of it from the situational context or from the previous discourse. For example, the departure place may be inferred, because it is the arrival place of the previous section. The means to achieve this goal are utterances. Generating a semantic structure for an utterance may be considered as performing a speech act that alters the user's state of knowledge (Cohen and Perrault 1979). According to this view, the use of a certain utterance is limited by constraints referring to the user's current state of knowledge, and the form and content of previous utterances. Within the boundaries of these constraints, planning is assumed to be a dynamic process directed by random choices. As a result, the output of the planner will vary considerably from one run to another. Thus, the text planner is not designed to generate a plan that will eventually transfer the information to the user optimally, but instead to generate as many plans as possible, which nevertheless transfer the information in an acceptable way.

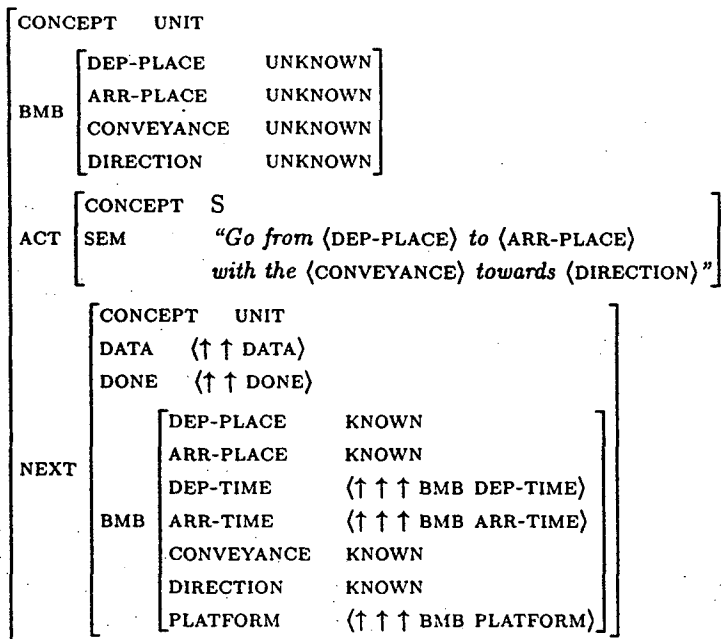


Figure 3: One of the grammar alternatives for a unit

The text planner is implemented as a Functional Unification Grammar (Kay 1984) in FUF (Elhadad 1993). The grammar is a feature description that consists of a number of alternatives, most of which represent an utterance with its constraints on application, its semantic structure and its effect on the user's knowledge. The process of text planning is a step-wise unification of the input with the grammar. The control mechanism of FUF traverses all concepts in the input structure (i.e. sub-FD's that contain the attribute CONCEPT), unifying them with suitable alternatives of the grammar. During this process, the input structure is enriched with new concepts, semantic structures and updates of the user's knowledge state.

We will trace this process on the basis of a simplified example. Suppose we take the FD in Figure 1 as input. Each SECTION concept in the input is unified with a corresponding grammar alternative. The grammar alternative for SECTION (see Figure 2) adds a feature UNITS that is used to store a number of nodes of type UNIT, corresponding to the utterances that together describe a section. A section typically contains between two and six units. A unit has a feature BMB, shorthand for 'belief-mutual-belief', which represents the text planner's belief about the current knowledge shared with the user. The alternative for SECTION initializes the knowledge state for its first unit: it is assumed that initially all information is unknown. The remaining features will be explained later on.

The grammar contains many different alternatives for UNIT, of which the one in Figure 3 is an example. The value of the BMB is best viewed as a condition on the applicability of this alternative. For the current example, it states that the departure place, departure time and conveyance must be unknown. Notice that, due to the nature of unification, the condition is indifferent with respect to the status of other data; they can be either known or unknown. If the condition succeeds, the speech act under ACT can be performed, which amounts to sending a semantic structure to the surface generator. The string template shown as the value of ACT is for expository reasons only; the value is actually an FD that is the semantic structure for an utterance. Semantic structures will be discussed later on. The slots in the template are filled by reference to the relevant values under DATA, which is the reason why this attribute is shared

CONCEPT		UNIT
	DEP-PLACE	KNOWN
	DEP-TIME	KNOWN
	ARR-PLACE	KNOWN
BMB	ARR-TIME	KNOWN
	CONVEYANCE	KNOWN
	DIRECTION	KNOWN
	PLATFORM	KNOWN
DONE	TRUE	

Figure 4: The special unit that states the termination condition

between a unit and a section (cf. Figure 2) and between units (cf. Figure 3).<sup>2</sup> Now performing a speech act alters the knowledge state, which is modeled by the fact that in the subsequent unit the values of the attributes DEP-PLACE, ARR-PLACE, CONVEYANCE, and DIRECTION become known. The state of the other data is shared with the previous BMB, implying that their status remains unaffected by the current speech act.

The expansion of a unit into a speech act and a next unit is a recursive process. It continues until BMB reaches the point where all data have become known. This termination condition is modeled by a special unit that has neither a speech act nor a NEXT attribute; see Figure 4. It does, however, provide the attribute DONE with its value TRUE, and because this value is shared between subsequent units as well as between a section and its first unit, it means that in the section node the attribute DONE becomes TRUE too. This in turn, triggers the alternatives<sup>3</sup> in Figure 2, which had been frozen by means of the special option *:wait* until the feature DONE had received a value. FUF tries the alternatives in the order they are given in the grammar. The first alternative succeeds if no more sections are given in the input, i.e. this was the last section of the route. Otherwise, the second alternative is taken, which forces processing of the next section.<sup>4</sup>

The important thing to notice is that when the next unit must be added, there are in general multiple units whose conditions are compatible with the current knowledge state. At such points, the random choice of a unit introduces the variation that was sought after. However, not every choice will lead to a solution, causing FUF to backtrack and revise its choice of units. Thus, the text planner can actually be considered a *planner* in the AI sense of the word as a program that traverses a search space (a network of connected units) for a path (a sequence of units with associated speech acts and knowledge updates) that satisfies its goal (a state where the planner believes that all data is shared with the user).

Figure 5 shows an example of a part of the output of the text planner based on the input in

<sup>2</sup>The fact that two attributes share the same value is expressed by means of *path*. For instance, the path  $\langle \uparrow \uparrow \text{DATA} \rangle$  that is the value of DATA means: go up two levels (i.e. skipping the attributes DATA and UNITS) and from there follow the attribute DATA to arrive at the intended value. This value is not present yet in the grammar alternative of Figure 2, but will be present in the input structure it is unified with.

<sup>3</sup>Alternatives (disjunctions) in the grammar are indicated by braces.

<sup>4</sup>The *:wait* option forces goal-freezing and is one of the ways in which FUF extends the FUG formalism it is based on. Another extension is exemplified by the special attribute CSET. By default, the unifier identifies constituents (i.e. sub-FD's in the input that need to be unified with the grammar) by the presence of a special attribute (CONCEPT in our case) and traverses these constituents in a top-down breadth-first manner. The CSET attribute enables the grammar writer to overrule this default and explicitly specify the constituents. This option is used to force processing of all units (by  $\text{CSET} == (\text{UNITS})$ ) before the next section is processed (by  $\text{CSET} + (\text{NEXT})$ ). See (Elhadad 1993).

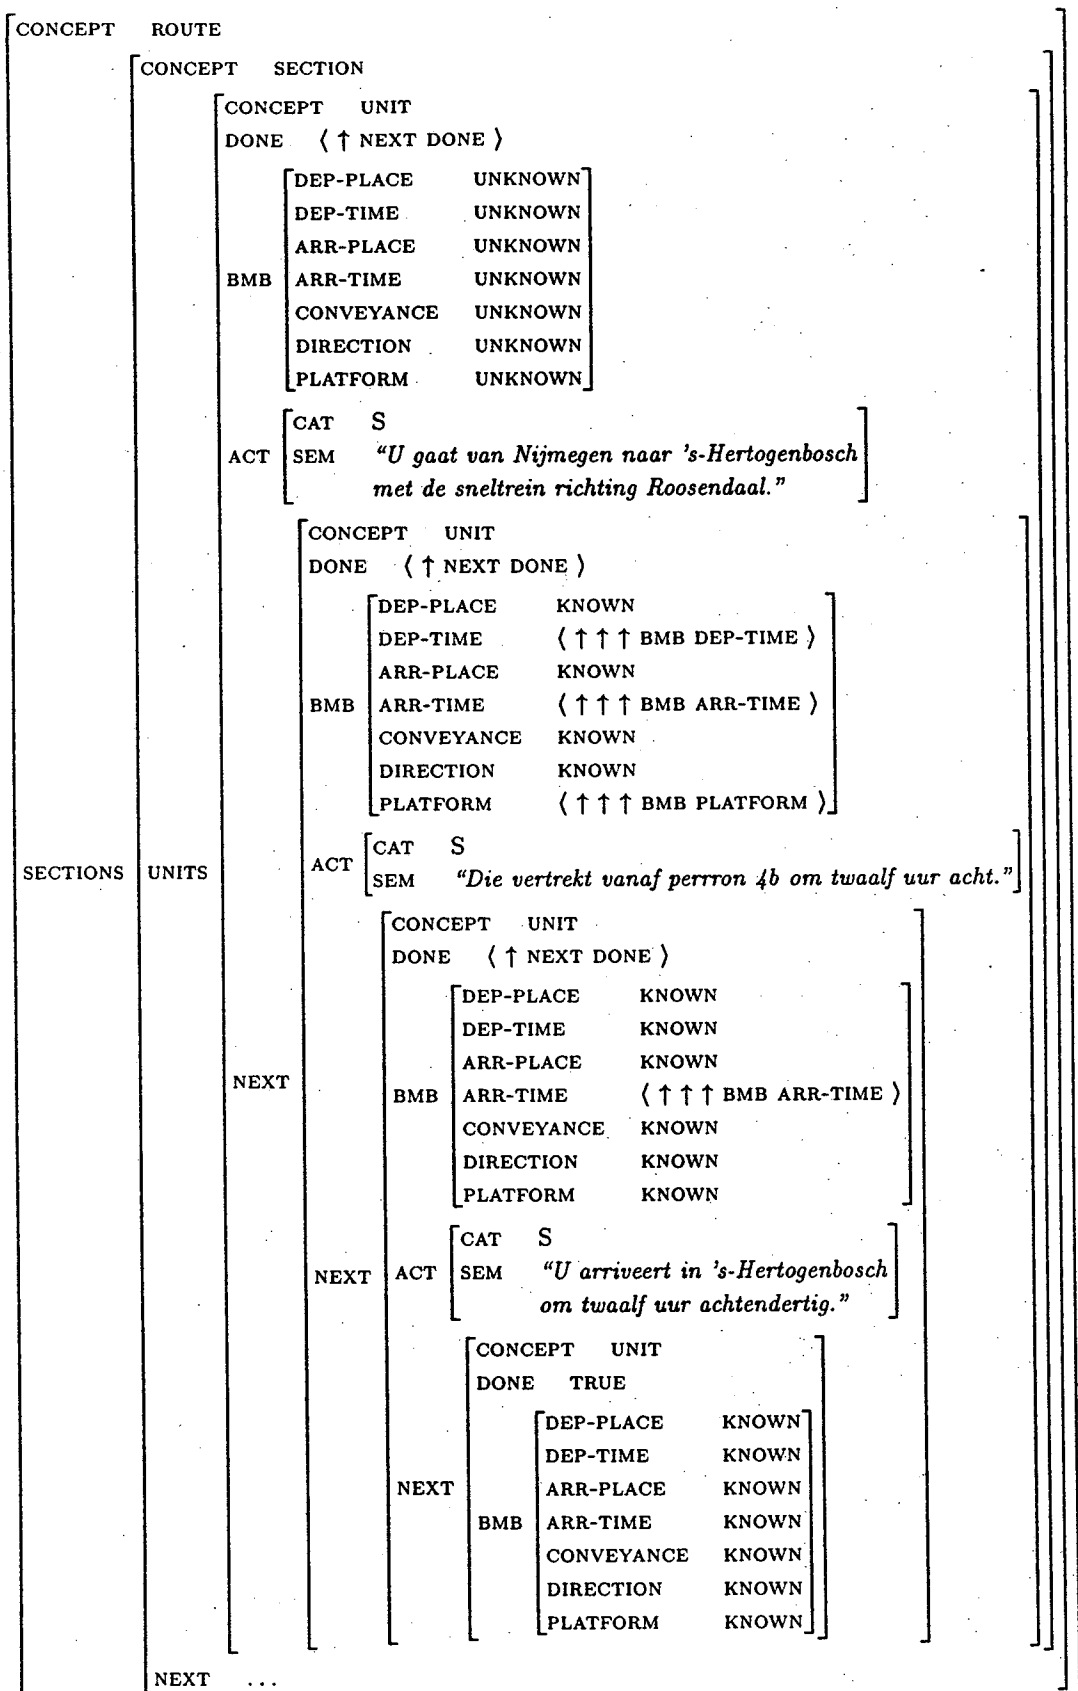


Figure 5: A part of the output that corresponds to the description of one section with three utterances. See example (3) for a gloss of the Dutch sentences.

Figure 1<sup>5</sup>. Notice that in the third unit the arrival place is repeated, although it was already mentioned in the first unit. This is possible because the grammar alternative for the third unit requires the departure time to be unknown, but does not constrain the value for arrival place. Therefore, it can be applied to introduce the arrival time only, or to introduce the arrival place as well. Either way, the arrival place is known after application of the unit. However, every units, with the exception of the termination unit, requires at least one piece of data to be unknown, since otherwise its application would be superfluous.

The planning grammar presented so far is simplified; the one actually used has a number of extensions. For instance, the assumption that some information is optional is modeled by relaxing the termination condition. That is, if the feature [PLATFORM KNOWN] is removed from the FD in Figure 4, then processing of a section may finish without making mention of the platform. Furthermore, the assumption that the place of departure is inferable, since it is the arrival place of the previous section, is implemented by forcing the departure place to be known in the first unit of a non-initial section. Two other extensions, for generating anaphoric expressions and discourse markers, will be discussed next.

### 2.3 Semantic structures

As mentioned earlier, the value of an ACT attribute is not a string template, but an FD that is the semantic structure for an utterance. An example is given in Figure 6.<sup>6</sup> It is passed on to a surface generator for Dutch that is similar to the SURGE surface generator for English (Elhadad and Robin 1996).<sup>7</sup> Notice how the lemmas for participants and circumstances are instantiated by means of paths that refer to the relevant values within the unit's DATA feature.<sup>8</sup>

Figure 6 also illustrates the distribution of focus. A constituent that is focused is presented as important to the listener (as opposed to unfocused material that is presented as less important to the listener). In general, information of which the speaker assumes that the listener is unfamiliar with is unfocused, and vice versa.<sup>9</sup> The distinction has repercussions for both syntactic and prosodic realization. Focus affects the syntactic structure, because it is used by the surface generator to determine the word order of an utterance. In particular, it will strive for a canonical word order with unfocused material at the start of the utterance and focused material at the end. Focus affects the prosodic structure, because focused material will be marked by at least one pitch accent. For current purposes, this means that checking the value of BMB provides a convenient way to determine if something is focused or not. This check is implemented as the option between parentheses in Figure 6. It states that if the arrival place is known, then its realization must be unfocused. However, if the arrival place is unknown, the option fails and the value for FOCUS is left unspecified. This interacts with the default assumption about focus made by the surface generator: content words are focused, while function words are unfocused. Hence, the text planner can limit itself to the exceptions, like the aforementioned case where the departure place is realized as a content word, but is nonetheless unfocused. Likewise, there is no need to explicitly specify that the Instrument is focused, or the Agent is unfocused.

In addition to the distribution of focus, the text planner is also responsible for generating

<sup>5</sup>The features DATA, CSET, FC, as well as the second section, were left out to save space.

<sup>6</sup><↑<sup>7</sup> DATA ARR-PLACE> is an abbreviation of <↑↑↑↑↑↑↑↑ DATA ARR-PLACE>

<sup>7</sup>This generator, called SEM2SYN, is a reusable surface generator for Dutch implemented in FUF (Marsi 1998). Its use is not limited to the present domain of travel descriptions. It has also been used to generate botanical descriptions of plants.

<sup>8</sup>At present, the text planner performs lexical choice, and is therefore responsible for variation at lexical level. This is not a not the only option however, since lexical choice might as well be performed in a separate module.

<sup>9</sup>The focused versus unfocused distinction does not always coincide with the known versus unknown distinction. For example, old or known information may be focused to obtain a contrastive effect.

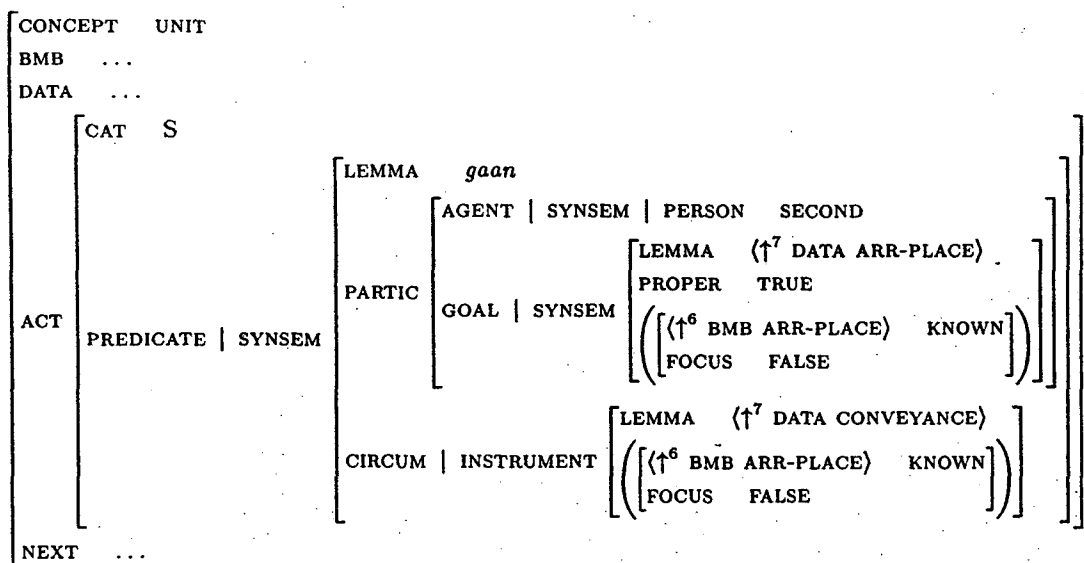


Figure 6: A unit containing the semantic structure for the utterance *U gaat naar <arr-place> met de <conveyance>*. ‘You go to <arr-place> with the <conveyance>.’

anaphoric expressions. The range of possible anaphoric expressions within the present domain is quite small. First, the listener is situationally evoked and is always referred to by a personal pronoun. Second, the conveyance may be referred to by a relative pronoun if it has been mentioned before. Third, a departure place, arrival place, direction or platform may be referred to by a locative anaphoric adverb. The latter type of reference is less trivial, because its use is restricted by word order. A case in point is (1) versus (2). The anaphoric expression *daar* (‘there’) is most naturally interpreted as referring to the place that was most recently mentioned. This leads to the intended interpretation (i.e. the departure place) in (1-b), but to a confusing or even unintended interpretation (i.e. the direction of the conveyance) in (2-b). Thus, in order to generate adequate anaphoric expressions of place, the text planner must keep track of the most recently mentioned place.

- (1) a. *U neemt de sneltrein richting Roosendaal in Nijmegen<sub>i</sub>.*  
you take the train towards Roosendaal in Nijmegen
- b. *Daar<sub>i</sub> vertrekt u om 12:08 van perron 4b.*  
there leave you at 12:08 from platform 4b
- (2) a. *U neemt in Nijmegen de sneltrein richting Roosendaal<sub>i</sub>.*  
you take in Nijmegen the train towards Roosendaal.
- b. *\*Daar<sub>i</sub> vertrekt u om 12:08 van perron 4b.*  
there leave you at 12:08 from platform 4b

This is implemented by means of a feature FC<sup>10</sup> that tells a unit what the most recently mentioned items of type HUMAN, PLACE and OBJECT are. This way, a unit can consult the content of FC to decide if an anaphoric expression can be used in its accompanying utterance. Depending on the content and word order of its utterance, a unit projects similar information to the FC

<sup>10</sup>FC stands for ‘forward centers’, because its use shows some resemblance to the notion of a set of forward centers in centering theory (Grosz, Joshi, and Weinstein 1995). However, the text planner is certainly not meant to be an implementation of centering theory.



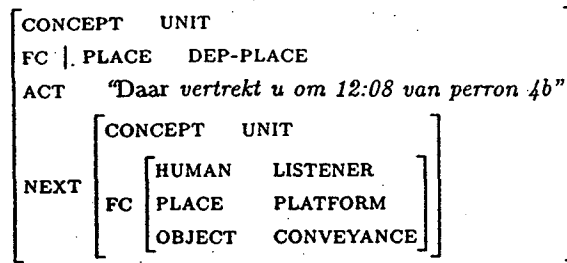


Figure 7: Example of the use of FC. See example (1a) for a gloss of the Dutch sentence.

feature of the next unit. An example is given in Figure 7.

The difference between the features BMB and FC is that the former tells us whether something is already *known* (either because it was mentioned in one of the previous sentences or because it could be inferred), whereas the latter tells us whether something was *mentioned* in the *latest* utterance.

This approach poses an interesting question regarding word order: is word order determined by the text planner, by the surface generator, or perhaps by both? One point of view is that a semantic structure is passed on to the surface generator, which determines the word order, which in turn determines the FC a unit projects to the next unit. This assumes that there is feedback from surface generator to text planner and that generation proceeds 'depth-first' (i.e. plan the first utterance, realize the first utterance, plan the second utterance, realize the second utterance, etc.) An alternative point of view is that the semantic structure contains restrictions on word order (like 'mention the departure place last'), depending on the FC a unit projects to the next unit. This requires no feedback and assumes that the generation process is 'breadth-first' (i.e. planning all utterances before sending them to the surface generator). So far, we have adopted the latter approach, because it is less complicated, both in concept and in implementation.

Finally, the text planner also inserts discourse markers. For the time being, this is just a provisional solution to improve the quality of the output; the implementation is not based on any theory. Since the nature of the domain is a small narrative in which the sections are described in chronological order, temporal continuity markers are suitable in most cases. For example, the first unit of non-initial section may add a temporal continuation marker like *next*, *then*, *after that* etc. It would be interesting to explore the possibilities of a more principled account of discourse markers, e.g. by using rhetorical relations as in (Hovy 1991).

## 2.4 Final output

(3) gives an example of a travel description in Dutch, generated by the combination of text planner and the surface generator, and based on the input in Figure 1.

- (3) a. *U gaat van Nijmegen naar 's-Hertogenbosch met de sneltrein richting Roosendaal.*  
you go from Nijmegen to 's-Hertogenbosch with the express-train towards Roosendaal.  
'You take the Roosendaal express train to 's-Hertogenbosch.'
- b. *Die vertrekt van perron 4b om twaalf uur acht.*  
that leaves from platform 4b at twelve hour eight  
'Which leaves from platform 4b at 12.08'

- c. *U arriveert in 's-Hertogenbosch om twaalf uur achtendertig.*  
You arrive in 's-Hertogenbosch at twelve hour thirty-eight  
'Which gets you to 's-Hertogenbosch at 12.38'
- d. *Vervolgens neemt u daar de stoptrein richting Utrecht Centraal Station.*  
next take you there the local-train towards Utrecht Central Station  
'Next, take the local train to Utrecht Central Station.'
- e. *Die vertrekt in 's-Hertogenbosch van perron 3b om twaalf uur tweeveertig.*  
that leaves in 's-Hertogenbosch from platform 3b at twelve hour fortytwo  
'Which leaves in 's-Hertogenbosch from platform 3b at 12.42.'
- f. *Dan bent u in Geldermalsen om twaalf uur negenenvijftig.*  
then are you in Geldermalsen at twelve hour fiftynine  
'Which gets you to Geldermalsen at 12.59.'

### 3 Summary

We have described a simple method for text planning that is suitable to small domains like travel information. Our aim was to introduce maximal variation in the packaging of information and in the linear order of its presentation. To this end, it proved useful to view text planning as a goal-driven process, in which utterances are used to alter the knowledge state of the user, and their use is restricted only by constraints that refer to the user's knowledge state. This can be conveniently implemented as a Functional Unification Grammar. In addition, we showed how optional or inferable information can be accounted for, how focus is distributed, and how the generation of anaphoric expressions can be constrained by the form and content of the latest utterance. Future work may address the necessity of feedback from surface generator to text planner, and the incorporation of a more principled account of generating discourse markers. Furthermore, an evaluation of the output with real users would be desirable.

### References

- Cohen, P. and C. Perrault (1979). Elements of a plan-based theory of speech acts. *Cognitive Science* (3), 177-212.
- Elhadad, M. (1993). *FUF: The Universal Unifier - User Manual, version 5.2*. New York. Technical Report CUCS-038-91.
- Elhadad, M. and J. Robin (1996). An overview of SURGE: A reusable comprehensive syntactic realization component. Technical Report 96-03, Ben Gurion University, Department of computer Science, Beer Sheva, Israel.
- Grosz, B. J., A. K. Joshi, and S. Weinstein (1995). Centering: a framework for modeling the local coherence of discourse. *Computational Linguistics* 21(2), 203-225.
- Hovy, E. H. (1991). Approaches to the planning of coherent text. In W. Swartout and W. Mann (Eds.), *Natural Language Generation in Artificial Intelligence and Computational Linguistics*, pp. 83-102. Boston: Kluwer Academics Publishers.
- Kay, M. (1984). Functional unification grammar: A formalism for machine translation. In *Proceedings of COLING-84*, pp. 75-78. ACL, Stanford University.
- Marsi, E. (1998). A reusable syntactic generator for Dutch. In P.-A. Coppen, H. V. Halteren, and L. Teunissen (Eds.), *Proceedings of CLIN97*. Nijmegen.