



FREIE UNIVERSITÄT BOZEN  
LIBERA UNIVERSITÀ DI BOLZANO  
FREE UNIVERSITY OF BOZEN - BOLZANO



Faculty of Computer Science, Free University of Bozen-Bolzano,  
Piazza Domenicani 3, 39100 Bolzano, Italy  
Tel: +39 04710 16000, fax: +39 04710 16009, <http://www.inf.unibz.it/krdb/>

*KRDB Research Centre Technical Report:*

# Natural Language Rendering of a Conjunctive Query

Paolo Dongilli

<b>Affiliation</b>	KRDB, Faculty of Computer Science, FUB Piazza Domenicani 3, 39100 Bolzano, Italy
<b>Corresponding author</b>	Paolo Dongilli <a href="mailto:dongilli@inf.unibz.it">dongilli@inf.unibz.it</a>
<b>Keywords</b>	conjunctive query, description logics, natural language generation, discourse planning, topological sorting, centering theory, discourse communication knowledge, sentence planning, sentence plan language, sentence aggregation, referring expression generation, linguistic realization, systemic-functional linguistics, systemic-functional grammar
<b>Number</b>	KRDB08-3
<b>Date</b>	16-06-2008
<b>URL</b>	<a href="http://www.inf.unibz.it/krdb/">http://www.inf.unibz.it/krdb/</a>

© **KRDB Research Centre**. This work may not be copied or reproduced in whole or part for any commercial purpose. Permission to copy in whole or part without payment of fee is granted for non-profit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of the KRDB Research Centre, Free University of Bozen-Bolzano, Italy; an acknowledgement of the authors and individual contributors to the work; all applicable portions of this copyright notice. Copying, reproducing, or republishing for any other purpose shall require a licence with payment of fee to the KRDB Research Centre.

# Natural Language Rendering of a Conjunctive Query

Paolo Dongilli

KRDB Center – Faculty of Computer Science  
Free University of Bozen-Bolzano, Italy  
dongilli@inf.unibz.it

June 16, 2008

## Abstract

The setting of this work is the attempt of building a bridge between an intelligent query interface we developed and state-of-the-art natural language generation (NLG) technologies. We built a query interface that allows the user to formulate a query over a knowledge base (KB) represented by a logic-based domain ontology. This report tackles each one of the main tasks in natural language generation (NLG), namely *text planning* (content determination, discourse planning), *sentence planning* (lexicalization, sentence aggregation, referring expression generation), and *linguistic realization* (syntactic and morphological realization, orthographic realization), presenting a pipeline of steps which build up our own NLG architecture, able to map a conjunctive query (over a given domain ontology) into its corresponding textual form.

## Contents

<b>1</b>	<b>Introduction</b>	<b>6</b>
<b>2</b>	<b>Conjunctive Queries over a Knowledge Base</b>	<b>7</b>
<b>3</b>	<b>Text Planning</b>	<b>9</b>
3.1	Content determination . . . . .	9
3.2	Discourse planning . . . . .	9
3.2.1	Centering-Theory-based planning . . . . .	13
3.2.2	Minimal conceptual distance . . . . .	16
3.2.3	Hybrid approach . . . . .	17
3.2.4	User-driven planning . . . . .	18
3.2.5	Depth-first planning . . . . .	20
3.2.6	Relation-priority depth-first planning . . . . .	20
3.3	Summary . . . . .	22

<b>4</b>	<b>Sentence Planning</b>	<b>23</b>
4.1	Sentence aggregation . . . . .	23
4.1.1	Types of Aggregation . . . . .	24
4.1.2	Aggregation Template Structures . . . . .	27
4.1.3	Best template structure matching . . . . .	30
4.1.4	Algorithms . . . . .	34
4.2	Referring expressions generation . . . . .	36
4.3	Generation of a Sentence Plan in SPL . . . . .	42
<b>5</b>	<b>Linguistic Realization</b>	<b>58</b>
5.1	Approaches to LR . . . . .	58
5.1.1	Canned Text Systems . . . . .	58
5.1.2	Template-based Systems . . . . .	59
5.1.3	Phrase-based Systems . . . . .	59
5.1.4	Feature-based Systems . . . . .	59
5.2	The Linguistic Realizer of Choice . . . . .	59
<b>6</b>	<b>Linguistic Realization with Systemic Functional Grammar</b>	<b>60</b>
6.1	Systemic Functional Grammar . . . . .	61
6.1.1	History . . . . .	61
6.1.2	Theory . . . . .	61
6.2	The Nigel systemic grammar of English . . . . .	64
6.3	The Upper Model . . . . .	66
6.4	Input specification: the Sentence Plan Language . . . . .	67
6.5	The KPML System . . . . .	69
6.5.1	The KPML Generation Process . . . . .	69
6.5.2	KPML Input Resources . . . . .	71
<b>7</b>	<b>Conclusions</b>	<b>76</b>

## List of Algorithms

1	Generation of a topological sort . . . . .	11
2	Generation of all topological sorts . . . . .	12
3	Generation of the best topological sorts (Centering Theory) . . .	15
4	Generation of the best topological sorts (Hybrid approach #1 (CT-mCD)) . . . . .	17
5	Generation of one of the best topological sorts (Hybrid approach #2 (mCD-CT)) . . . . .	18
6	Generation of a topological ordering using depth-first search . .	20
7	Linearization of a query tree . . . . .	34
8	Calculation of the best covering match . . . . .	35
9	Generation of appropriate referring expressions for each entity present in a given text plan . . . . .	40

## List of Figures

1	Example of query graph. . . . .	8
2	A query tree . . . . .	10

3	A discourse tree . . . . .	10
4	Best and worst cases for topological sorting . . . . .	12
5	Adding a new relation to the query tree of Fig. 2 . . . . .	19
6	Discourse tree derived from the query tree of Fig. 5 . . . . .	19
7	Query tree with ordering annotations attached to relations . . . . .	21
8	Discourse tree with ordering annotation attached to discourse units . . . . .	21
9	A query tree example. . . . .	24
10	A query tree waiting to be linearized . . . . .	31
11	Query with abstract role ( <i>make</i> ) that is rendered as substantive. . . . .	41
12	Example of system network fragment . . . . .	62
13	Example of system network fragment with realization statements . . . . .	64
14	Simplified example of metafunctional layering . . . . .	64
15	Example of system network with choosers . . . . .	66
16	Excerpt of the Penman Upper Model taxonomy . . . . .	67
17	KPML Penman-style generation architecture (based on [Bateman, 1997b]) . . . . .	70
18	Excerpt of LOOM ontology from the automotive domain . . . . .	72
19	Lexical items from the automotive domain . . . . .	75

## List of Tables

1	Aggregation template structures . . . . .	27
2	Aggregation examples . . . . .	29
3	Linearized templates . . . . .	31
4	Examples of usage of referring expressions . . . . .	38
5	Complete set of singular pronouns used . . . . .	42
6	Main systems in sFG . . . . .	63
7	Realization statements used in sFG . . . . .	64
8	Mapping of domain ontology entities and subordination to UM entities . . . . .	72
9	Functional regions in the Nigel Grammar for English . . . . .	73

# 1 Introduction

Communication by means of natural language involves two fundamental skills: Producing text and understanding it. These tasks are the subject of study of two big areas of research in computational linguistics, which are *natural language generation* and *natural language understanding* respectively; the former will be dealt with hereinafter.

Natural language generation (NLG) is seen in general as the sequence of operations needed to map information from some *non-linguistic* (e.g. raw data) into *linguistic* form (either oral or written). These operations are not at all straightforward, because the task of bridging the gap between non-linguistic and linguistic representations requires several non-trivial decisions or choices which include *content determination*, choice of *rhetorical structures* at various levels (text, paragraph, sentence), choice of *words* and *syntactic structures*, and finally the determination of the *text layout* (or acoustic patterns if we intend to generate spoken text). One of the main challenges of NLG is devising modular architectures able to make the previous choices coexist. At least three kinds of expertise are needed: *application domain knowledge*, *knowledge of the language* (grammar, lexicon, and semantics), and *strategic rhetorical knowledge* (i.e. how to achieve communicative goals, text types, style).

NLG system architectures need to include various levels of planning and merging of information in a way that generated text looks natural and not repetitive. Typical tasks we find are [Reiter and Dale, 2000]:

- **Text planning**
  - **Content determination:** Determination of the salient features that are worth being said.
  - **Discourse planning:** Overall organization of the information to convey.
- **Sentence planning**
  - **Lexicalization:** Putting words to the concepts.
  - **Sentence aggregation:** Merging of similar sentences to improve readability and naturalness. For example, the sentences “The car is equipped with a diesel engine” and “The engine’s power is 140 HP” can be aggregated to form “The car is equipped with a diesel engine whose power is 140 HP”.
  - **Referring expression generation:** Linking words in the sentences by introducing pronouns and other types of means of reference.
- **Linguistic realization**
  - **Syntactic and morphological realization:** This stage is the inverse of parsing: given all the information collected above, syntactic and morphological rules are applied to produce the surface string.
  - **Orthographic realization:** Matters like casing, punctuation, and formatting are resolved

In the coming sections we will tackle each one of the mentioned tasks, describing a pipeline of steps which will build up our own generation architecture, able to map a conjunctive query (over a given domain ontology) into its corresponding textual form. We first start by defining what our input is, i.e. a conjunctive query over a knowledge base (KB).

## 2 Conjunctive Queries over a Knowledge Base

The type of queries we present in this section are *conjunctive queries* (CQs) over a given Description Logics [Baader *et al.*, 2007] knowledge base. Their definition will be given along with the definitions of *query answering* and *query graphs*.

**Definition 1** (Querying a KB). Querying a (Description Logics) KB means verifying whether a given statement  $q$  (the query) is a logical consequence of the knowledge base (i.e.  $K \models q$ ).

Given a KB  $K$ , it is possible to query it using *conjunctive queries* which are defined as follows:

**Definition 2** (Conjunctive Queries). A *conjunctive query* can be represented as  $q(\vec{x}) \leftarrow \text{conj}(\vec{x}, \vec{y})$ , where  $\vec{x}$  is the vector of so called *distinguished variables* that will be bound to individuals (single objects) of the knowledge base used to answer the query;  $\vec{y}$  is the vector of non-distinguished variables (existentially quantified variables).  $\text{conj}(\vec{x}, \vec{y})$  is a conjunction of terms of the form  $v_1 : C$ ,  $\langle v_2, v_3 \rangle : R$ , or  $v_4 : P$  where  $C$  is a concept name,  $R$  is a role name (*abstract role* or *relation* if  $v_3$  refers to an abstract concept, *concrete role* or *attribute* if  $v_3$  refers to a concrete domain),  $P$  is a simple or complex predicate over a given concrete domain  $\mathcal{D}$  and  $v_1, v_2, v_3, v_4$  are variables from  $\vec{x}$  or  $\vec{y}$ .

**Example 1.** A query example taken from the automotive domain could be the following:

$$\begin{aligned} q(\underline{x}_1, \underline{x}_6, \underline{x}_7, \underline{x}_8) \leftarrow & \underline{x}_1 : \text{Car} \sqcap \langle \underline{x}_1, \underline{x}_2 \rangle : \text{equippedWith} \sqcap \\ & \underline{x}_2 : \text{Engine} \sqcap \langle \underline{x}_2, \underline{x}_3 \rangle : \text{runningOn} \sqcap \underline{x}_3 : \text{Diesel} \sqcap \\ \langle \underline{x}_1, \underline{x}_4 \rangle : \text{madeBy} \sqcap & \underline{x}_4 : \text{LandRover} \sqcap \langle \underline{x}_1, \underline{x}_5 \rangle : \text{soldBy} \sqcap \\ \underline{x}_5 : \text{CarDealer} \sqcap & \langle \underline{x}_2, \underline{x}_6 \rangle : \text{displSizeCC} \sqcap \underline{x}_6 : \text{integer} \sqcap \\ \langle \underline{x}_5, \underline{x}_7 \rangle : \text{locInCity} \sqcap & \underline{x}_7 : \text{City} \sqcap \langle \underline{x}_5, \underline{x}_8 \rangle : \text{name} \sqcap \\ \underline{x}_8 : \text{string} \sqcap & \langle \underline{x}_7, \underline{x}_9 \rangle : \text{locInCountry} \sqcap \underline{x}_9 : \text{Italy} \sqcap \\ & \langle \underline{x}_7, \underline{x}_{10} \rangle : \text{locInProvince} \sqcap \underline{x}_{10} : \text{Trento} \end{aligned}$$

*Underlined variables represent distinguished variables. Uppercase names are abstract concepts, and lowercase names are relation names; string and integer represent instead two concrete domains.*

*The query above is for “a car equipped with an engine running on Diesel. The car is a Land Rover and it is sold by a car dealer located in Italy in the province of Trento. The user wants to know the cars available, the engine displacement size, the car dealer’s name, and the city.”*

We explain now what it means answering a conjunctive query followed by the definition of query graph.

**Definition 3** (Conjunctive Query Answering). Given a query  $q(\vec{x})$  where  $\vec{x}$  are distinguished variables, and a KB  $K$ , answering  $q(\vec{x})$  means returning all tuples  $\vec{t}$  that substituted to  $\vec{x}$  are such that  $K \models q(\vec{t})$ .

**Definition 4** (Query Graphs). A conjunctive query  $q$  can be represented by means of a directed labelled graph  $G(q) := \langle V, E \rangle$  where  $V$  represents a set of vertices and  $E$  a set of edges.  $V$  is the union of the elements in  $\vec{x}$ , and  $\vec{y}$ ;  $E$  is made up of all pairs  $\langle v_1, v_2 \rangle$  where  $v_1, v_2 \in V$  and  $\langle v_1, v_2 \rangle : R$  is a term in  $q$ . A node  $v \in V$  is labelled with a concept  $C_1 \sqcap \dots \sqcap C_n$  such that for every  $C_i$ ,  $v : C_i$  is a term of  $q$ . Optionally if node  $v$  is labelled with a concrete domain, it can also have an additional label containing the name of a (simple or complex) predicate. Every edge  $e \in E$  is labelled with a set of role names  $\{R \mid \langle v_1, v_2 \rangle : R \text{ is a term in } q\}$ .  $q$  is an *acyclic conjunctive query* if  $G(q)$  is not cyclic.

For a better graph readability, a node  $v$  is represented with a rectangle when it refers to an abstract concept, and with an oval when it refers to a concrete domain.

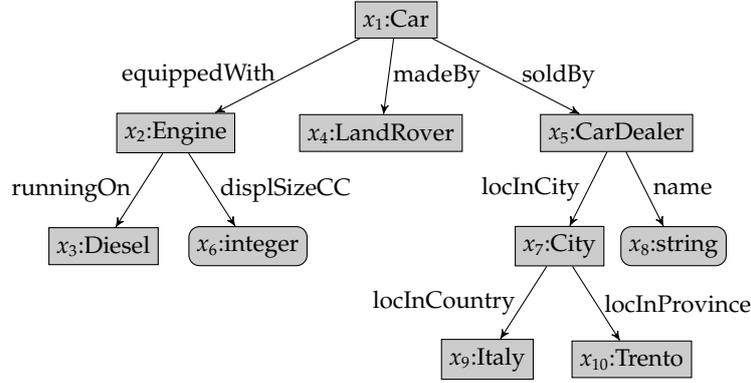


Figure 1: Example of query graph.

Hereinafter we will use the terms *query* or *conjunctive query* interchangeably referring to acyclic conjunctive queries, and the corresponding query graph will be called *query tree*. Figure 1 shows the query tree corresponding to the query in Example 1.

It is possible to describe our query from the point of view of a single variable (the *focus*). This new query is called *focused query*; it is equal to the original one, with the exception that the only distinguished variable of the focused query is the variable representing the focus.

In the example above we could have a query  $q^{x_1}$  focused on  $x_1$  (a Car), or a query  $q^{x_5}$  focused on  $x_5$  (CarDealer). This transformation turns out to be very useful, because given the restriction to tree-shaped conjunctive query expressions, together with the availability of inverse roles, a focused query corresponds to a concept expression [Horrocks and Tessaris, 2002], on which we can perform operations by means of standard DL reasoning services (mainly satisfiability checking and classification), and through these we can drive the query interfaces we implemented (see [Dongilli *et al.*, 2004; Zorzi *et al.*, 2007]).

With these premises, we describe now our approach for the generation of a suitable text plan from a given query which represents our input. This will

be followed by the description of the sentence planning phase and finally the linguistic realization phase.

### 3 Text Planning

Thinking what to write or say, and organizing the constituents of our idea in one of the possible manners that once verbalized will best convey our thought is what we could define as the *text planning* capability of a human being. We want to mimic this human behavior with the first module of our generation system, where the content (a query), determined by the user by means of an intelligent query composition interface, is internally reorganized in order to obtain the possibly most coherent sequence of its constituents. Since the query is isomorphic to a tree, the job of the discourse planner is to find the best topological sorting according to some objective function. In Section 3.2 we present and compare six discourse planning strategies indicating one of them as the best suited for this task. These results were also published in [Dongilli, 2007a] and [Dongilli, 2007b].

#### 3.1 Content determination

In our specific context, the content is represented by the query formulated by the user. Given the specific domain of interest chosen (read *ontology*), the query built using the ontology's (unary and binary) predicates corresponds to a simple or complex concept (namely a conjunction of predicates) whose constituents need to be organized and explained in a coherent discourse using one or more natural language sentences.

#### 3.2 Discourse planning

Planning a discourse means in general finding the best way of representing an idea in an organized, specific, and coherent manner. In our case the idea is a query, a complex concept that the user is thinking and building. We start by defining the main components of a discourse that are called *discourse units*.

**Definition 5** (Discourse unit). A discourse unit  $u_i(c_j, c_k)$  is the atomic component of a discourse. In our setting a discourse unit is represented by a *role*  $r_l$  between two entities  $c_j$  and  $c_k$  (using the terminology introduced by def. 2,  $r_l$  can be either a relation or an attribute having  $c_j$  and  $c_k$  as domain and range respectively;  $c_j$  is always a concept, and  $c_k$  is a concept if  $r_l$  is a relation, or a predicate over a concrete domain if  $r_l$  is an attribute).

A discourse unit, once verbalized, can be seen as a stand-alone sentence, or as a clause in a longer sentence.

We start with a query tree and we map it into another tree we call *discourse tree*, created by collapsing a role between two concepts into a single node.

**Definition 6** (Discourse tree). A discourse tree is a directed tree whose nodes are discourse units. The nodes are tagged with the domain and range entities of the corresponding role. The edges connect two nodes where the second entity of the start node and the first entity of the end node are the same. The root node

of the discourse tree is an additional node which introduces the main concept ( $c_1$ ) the user is looking for, i.e. the root concept of the query tree. The first entity ( $c_0$ ) of this node is a new entity that will be verbalized as the subject of this first discourse unit.

Figures 2 and 3 show a starting query tree and the derived discourse tree. The sequence order of indexes assigned to concepts and relations in the query tree, and as consequence the indexes assigned to discourse units, respects the order of insertion followed by the user while creating the query.

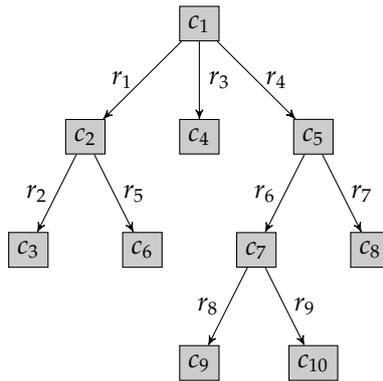


Figure 2: A query tree

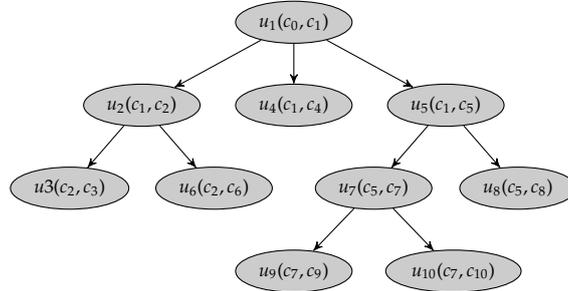


Figure 3: A discourse tree

Our starting point for discourse planning is the generated discourse tree which is a directed tree as mentioned in def. 6. The problem of finding a linear sequence of the discourse units in a discourse tree can be translated into a problem which in graph theory is called *topological sorting*.

**Definition 7** (Topological sort). A *topological sort* of a directed acyclic graph (DAG)  $G(V, E)$  is a linear ordering of its nodes  $V$  which is compatible with the partial order  $R$  induced on the nodes by the set of directed edges  $E$ , where  $x$  comes before  $y$  ( $xRy$ ) if there's a directed path from  $x$  to  $y$  in the DAG (where  $x, y \in V$ ).

In other terms, topological sorting is a way to extend a *partial order* relation into a *total order*. We can state that every DAG has at least one topological sort, because of the following

**Theorem 1.** *Every partial order can be extended to a total order. That is: Suppose  $\rightarrow$  is a partial order on a set  $X$ . Then there exists a total order  $\Rightarrow$  on  $X$  that extends  $\rightarrow$  as a relation: If  $x, y \in X$  and  $x \rightarrow y$ , then  $x \Rightarrow y$ .*

Typical algorithms for topological sorting have running time linear in the number of nodes plus the number of edges ( $\Theta(|V| + |E|)$ ). Since in our setting we are working with a DAG where  $|E| = |V| - 1$ , the complexity is  $\Theta(|V|)$ . A possible algorithm is the following:

---

**Algorithm 1** Generation of a topological sort

---

```

Q ← Set of all nodes with no incoming edges
while Q is not empty do
  remove a node n from Q
  output n
  for all nodes m with an edge e from n to m do
    remove edge e from the graph
    if m has no other incoming edges then
      insert m into Q
    end if
  end for
end while
if graph has edges then
  output error (the graph has a cycle)
end if

```

---

If we want to find all possible topological sorts, algorithm 1 needs the following modifications that lead to Algorithm 2:

- the *while* loop must be implemented by a recursive function;
- set  $Q$  has to be global to the recursive function;
- an array must be defined to hold the current ranking and that must be output when  $Q$  is found to be empty;
- the currently removed item is to be kept locally;
- when returning from the recursive call, the current item has to be put back into  $Q$  and another node must be picked from  $Q$ .

We implemented this algorithm (see Algorithm 2) whose running time depends on the topology of the tree. It is easy to see that the number of topological sorts varies from 1 to  $|E|!$ ; these two extreme cases are shown in Figure 4. The former tree already represents a linear order, while the latter has  $n - 1$  possible topological sorts.

In our context, finding a topological sort of a discourse tree can be defined this way:

**Definition 8** (Topological sort of a discourse tree). Given a discourse tree with  $n$  discourse units  $u_1, u_2, \dots, u_n$  containing  $n + 1$  discourse entities  $c_0, c_1, c_2, \dots, c_n$ , a *topological sort* can be obtained with a permutation  $\pi$  of  $\{1, 2, \dots, n\}$ , where the sequence of discourse units  $(u_{\pi(1)}, u_{\pi(2)}, \dots, u_{\pi(n)})$  is compatible with the partial order induced by the discourse tree.

---

**Algorithm 2** Generation of all topological sorts

---

$Q \leftarrow$  Set of all nodes with no incoming edges  
 $D \leftarrow \emptyset$  {array containing temporary linear sort}  
call *CalculateSorts*()

**procedure** CALCULATESORTS()

**if**  $Q$  is not empty **then**

**for**  $i = 1$  to  $\text{size}(Q)$  **do**

    remove node  $n_i$  from  $Q$

    add  $n_i$  to  $D$

**for all** nodes  $m$  with an edge  $e$  from  $n_i$  to  $m$  **do**

      remove edge  $e$  from the graph

**if**  $m$  has no other incoming edges **then**

        insert  $m$  into  $Q$

**end if**

**end for**

    call *CalculateSorts*()

    restore  $n_i$  into  $Q$

    restore previously removed edges outgoing from  $n_i$

    remove last element from  $D$

**end for**

**else**

  output  $D$

**end if**

**if** graph has edges **then**

  output error (the graph has a cycle)

**end if**

**end procedure**

---

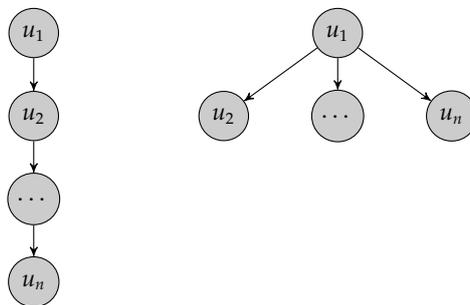


Figure 4: Best and worst cases for topological sorting

Hereinafter a generic topological sort will be associated and identified with a permutation  $\pi$ . Given all possible topological sorts of our discourse tree, we need now to find some constraints with the intent to keep only those orderings that maximize/minimize some objective function. By means of an objective function, the aim is to find some common properties of the best orderings, in a way to be able to infer an algorithm that is able to discern just one of the best topological sorts.

### 3.2.1 Centering-Theory-based planning

The constraints we use in this first attempt are borrowed from Centering Theory which gives us the means to find all possible sequences of discourse units that maximize coherence.

Centering theory (CT) finds its origins within the theory of discourse structure that was first developed by [Grosz and Sidner, 1986]. A draft manuscript describing the centering framework and the first theoretical claims appeared in 1986 [Grosz *et al.*, 1986], and the authors were then urged to publish a more detailed description which appeared in 1995 [Grosz *et al.*, 1995]. This, along with a previous contribution from [Brennan *et al.*, 1987], contains the main claims of this theory, which are:

1. for each discourse unit, there is exactly one entity which is the *center* of attention;
2. there is a preference for consecutive discourse units that keep the same entity as center, and for the most salient entity in a discourse unit to be realized as the center of the next utterance;
3. the center is the entity with the highest probability to be pronominalized.

The assumptions of CT are formalized in terms of  $C_f$ ,  $C_b$ , and  $C_p$ . Given two consecutive discourse units  $u_{\pi(i-1)}$  and  $u_{\pi(i)}$ ,

- $C_f(u_{\pi(i)})$  (*forward looking centers*) is a list of all discourse entities contained in  $u_i$ ;
- $C_b(u_{\pi(i)})$  (*backward looking center*) is the most highly *ranked* entity realized in  $u_{\pi(i-1)}$  which is also realized in  $u_{\pi(i)}$ ; If  $u_{\pi(i-1)}$  does not exist, there is no  $C_b(u_{\pi(i)})$ ;
- $C_p(u_{\pi(i)})$  (*preferred center*) is the highest *ranked* entity of  $u_{\pi(i)}$ .

[Brennan *et al.*, 1987] define *ranking* of an entity in a discourse unit as *the likelihood that it will be the primary focus of subsequent discourse*. It is more common now defining the rank in terms of grammatical roles (obliqueness), where subject > direct object > indirect object > others.

With the abovementioned parameters, we list now the following constraints, whose violations will build-up the cost function we are going to use.

**cohesion:**  $C_b(u_{\pi(i)}) = C_b(u_{\pi(i-1)})$  (checks if the center of the current discourse unit is the same as the preceding one);

**salience:**  $C_b(u_{\pi(i)}) = C_p(u_{\pi(i)})$  (checks if the center is realized as subject);

**cheapness:**  $C_b(u_{\pi(i)}) = C_p(u_{\pi(i-1)})$  (checks if the current center was a subject in the previous discourse unit);

**continuity:**  $C_f(u_{\pi(i)}) \cap C_f(u_{\pi(i-1)}) \neq \emptyset$  (checks whether (or not) there are entities in common between the previous and the current discourse unit).

We say that there is a *violation* to one of these constraints, if the corresponding condition does not hold. If there exists no  $C_b(u_{\pi(i)})$ , cohesion, salience, and cheapness are not violated. As a consequence, no violation is accounted for the first discourse unit  $u_{\pi(1)}$ .

The cost function we want to minimize in order to maximize local coherence is defined as follows:

**Definition 9** (Centering-theory-based cost function). Given the setting of definition 8, we define this cost function:

$$\phi_{CT}(\pi) = \sum_{i=1}^n [coh(u_{\pi(i)}) + sal(u_{\pi(i)}) + che(u_{\pi(i)}) + con(u_{\pi(i)})]$$

where:

$$\begin{aligned} coh(u_{\pi(i)}) &= \begin{cases} 0 & \text{if } i \in \{1, 2\} \text{ or} \\ & i > 2 \text{ and } C_b(u_{\pi(i)}) = C_b(u_{\pi(i-1)}) \\ K_{coh} & \text{if } i > 2 \text{ and } C_b(u_{\pi(i)}) \neq C_b(u_{\pi(i-1)}) \end{cases} \\ sal(u_{\pi(i)}) &= \begin{cases} 0 & \text{if } i = 1 \text{ or} \\ & i > 1 \text{ and } C_b(u_{\pi(i)}) = C_p(u_{\pi(i)}) \\ K_{sal} & \text{if } i > 1 \text{ and } C_b(u_{\pi(i)}) \neq C_p(u_{\pi(i)}) \end{cases} \\ che(u_{\pi(i)}) &= \begin{cases} 0 & \text{if } i = 1 \text{ or} \\ & i > 1 \text{ and } C_b(u_{\pi(i)}) = C_p(u_{\pi(i-1)}) \\ K_{che} & \text{if } i > 1 \text{ and } C_b(u_{\pi(i)}) \neq C_p(u_{\pi(i-1)}) \end{cases} \\ con(u_{\pi(i)}) &= \begin{cases} 0 & \text{if } i = 1 \text{ or} \\ & i > 1 \text{ and } C_f(u_{\pi(i)}) \cap C_f(u_{\pi(i-1)}) \neq \emptyset \\ K_{con} & \text{if } i > 1 \text{ and } C_f(u_{\pi(i)}) \cap C_f(u_{\pi(i-1)}) = \emptyset \end{cases} \end{aligned}$$

$K_{coh}$ ,  $K_{sal}$ ,  $K_{che}$ , and  $K_{con}$  represent the weights assigned to each constraint violation.

We can now use this cost function to discern, among all orderings, the ones that minimize violations to local coherence in terms of cohesion, salience, cheapness, and continuity, where the respective weights are assigned according to the proposal of [Kibble and Power, 2004]:  $K_{coh} = K_{sal} = K_{che} = 1$ , and  $K_{con} = 3$ . The assumption made is that for every discourse unit  $u_i(c_k, c_l)$ , the preferred center  $C_p(u_i)$  (subject) will always be the first discourse entity i.e.  $c_k$ .

We implemented algorithm 2, and conducted experiments over several tree topologies, isolating all sortings that minimized the given cost function. The results obtained, quite unexpected, are reported in the next section.

The first result we observed is that salience is never violated. This can be formalized in the following theorem:

**Theorem 2.** *Given the setting of Definition 8, none of the topological sorts generated from the discourse tree violates the salience constraint.*

*Proof.* In order to check salience in any discourse unit  $u_{\pi(i)}(c_l, c_m)$  of a generic topological sort  $\pi$ , we have to check that  $C_b(u_{\pi(i)}(c_l, c_m)) = C_p(u_{\pi(i)}(c_l, c_m))$  where  $C_p(u_{\pi(i)}(c_l, c_m)) = c_l$  as we assumed above. To identify the  $C_b$  we need the previous discourse unit  $u_{\pi(i-1)}(c_j, c_k)$ . If it does not exist, this means that  $u_{\pi(i)}$  is the first discourse unit ( $u_{\pi(1)}$ ) and there is no salience violation. If we have a previous utterance, we distinguish two cases:

1.  $\{c_l, c_m\} \cap \{c_j, c_k\} = \emptyset$ : this means that there is no  $C_b$  in unit  $u_{\pi(i)}$ , therefore salience is not violated;
2.  $\{c_l, c_m\} \cap \{c_j, c_k\} \neq \emptyset$ ; the units cannot have two discourse entities in common because this would mean that either the two entities are the same ( $c_j = c_l \wedge c_k = c_m$ ) or that we have a cycle in our tree ( $c_j = c_m \wedge c_k = c_l$ ) which is impossible. We can have only one entity in common, i.e. the following four cases:
  - (a)  $c_j = c_l$ : This is a valid case; it implies that  $C_b(u_{\pi(i)}(c_l, c_m)) = c_l$  which is equal to  $C_p(u_{\pi(i)}(c_l, c_m))$ . The salience constraint is attended.
  - (b)  $c_j = c_m$ : This case is not valid, because this would imply that  $u_{\pi(i-1)}$  should occur after before  $u_{\pi(i)}$ , contradicting our hypothesis.
  - (c)  $c_k = c_l$ : This is a valid case; it implies that  $C_b(u_{\pi(i)}(c_l, c_m)) = c_l$  which is equal to  $C_p(u_{\pi(i)}(c_l, c_m))$ . The salience constraint is attended.
  - (d)  $c_k = c_m$ : This is not a valid case, since it would imply coreference and therefore a cycle in our tree which is impossible.

□

The second result we obtained pertains a common property shown by all best topological sorts, i.e. the ones that minimize the cost function. This result is expressed in the following proposition.

**Proposition 1.** *Given a discourse tree with  $n$  discourse units  $u_1, u_2, \dots, u_n$  containing  $n + 1$  discourse entities  $c_0, c_1, c_2, \dots, c_n$ , the topological sorts that minimize the cost function expressed in Definition 9 are all and only the ones returned by Algorithm 3.*

---

**Algorithm 3** Generation of the best topological sorts (Centering Theory)

---

$D \leftarrow u_1$  {array  $D$  stores one by one all best linear sorts; here it is initialized with the first discourse unit}  
 $n \leftarrow$  number of nodes in tree  
 $count \leftarrow 1$  {current size of  $D$ }  
call *CalculateBestSorts*( $u_1$ )

**procedure** *CalculateBestSorts*( $u$ )

$L \leftarrow$  all children of  $u$

$L_d \leftarrow$  children of  $u$  having at least one descendant

**if**  $L_d \neq \emptyset$  **then**

$P_L \leftarrow$  list of all permutations of  $L$  where last node of each permutation is in  $L_d$

**else**

$P_L \leftarrow$  list of all permutations of  $L$

```

end if
for all  $p \in P_L$  do
  append array  $p$  to  $D$ 
  for  $i = \text{size}(p)$  downto 1 do
    call  $\text{CalculateBestSorts}(p_i)$ 
  end for
   $\text{count} \leftarrow \text{count} + \text{size}(p)$ 
  if  $\text{count} = n$  then
    output  $D$ 
  end if
  remove array  $p$  from  $D$ 
   $\text{count} \leftarrow \text{count} - \text{size}(p)$ 
end for
end procedure

```

---

In plain words, the best topological sorts are the ones for which every discourse unit is followed by its remaining siblings, where the last sibling must be one with descendants (to allow a continuity in the discourse); for each sibling then, starting from the last one, the list of its children (if any) is output. E.g. one of the best topological sorts from the discourse tree of Figure 3 is  $(u_1, u_2, u_4, u_5, u_8, u_7, u_9, u_{10}, u_3, u_6)$ .

### 3.2.2 Minimal conceptual distance

The second constraint we decided to experiment with, calculates the sum over each discourse entity of the distances among discourse units where each entity is referenced.

This optimization problem can be described as follows:

**Definition 10** (Conceptual distance minimization). Given  $n$  discourse units,  $(u_1, u_2, \dots, u_n)$  embedding  $n + 1$  discourse entities  $(c_0, c_1, c_2, \dots, c_n)$ ; given a permutation  $\pi$  of  $\{1, 2, \dots, n\}$  where the sequence of discourse units  $(u_{\pi(1)}, u_{\pi(2)}, \dots, u_{\pi(n)})$  is compatible with the partial order induced by the tree, we create a hash table  $H\pi$  where its keys correspond to the discourse entities  $(c_0, c_1, c_2, \dots, c_n)$ , and each value  $H\pi(c_i)$  is a sorted list of indexes taken from  $\{1, 2, \dots, n\}$  and referring to some positions of the permutation.

We want to

$$\min \sum_{i=0}^n \delta_{\pi}(c_i)$$

where

$$\delta_{\pi}(c_i) = \begin{cases} \sum_{k=1}^{|H_{\pi}(c_i)|-1} (H_{\pi}(c_i)[k+1] - H_{\pi}(c_i)[k]) & \text{if } |H_{\pi}(c_i)| > 1 \\ 0 & \text{if } |H_{\pi}(c_i)| = 1 \end{cases}$$

We were able to find a common property of all topological sorts that minimizes the newly introduced constraint; it is expressed by the following proposition:

**Proposition 2.** *Given a discourse tree with  $n$  discourse units  $u_1, u_2, \dots, u_n$  containing  $n + 1$  discourse entities  $c_0, c_1, c_2, \dots, c_n$ , the topological sorts that minimize the measure*

of conceptual distance as of Definition 10 are those that derive from a depth-first-like visit of the tree, where at any point, given an output node, the valid visits of its subtrees are all the ones where the biggest (in terms of number of nodes) subtree comes last with no other constraints.

In this case, one of the best topological sorts of the tree in Figure 3 is  $(u_1, u_4, u_2, u_3, u_6, u_5, u_8, u_7, u_9, u_{10})$ .

### 3.2.3 Hybrid approach

The next step was to run topological sorting using a hybrid approach with both of the two previous constraints. We ran several tests on different tree topologies.

We first applied the constraints based on centering theory (CT-approach) followed by the calculation of the minimal conceptual distance (mCD-approach), i.e. after minimizing the cost function of the CT-approach, we applied the mCD-approach on the best orderings. The best results coming out are obviously a subset of the orderings found by algorithm 3. The best topological sorts are the ones for which every discourse unit is followed by its remaining siblings, starting from the ones with no children (if any) and continuing with the siblings in decreasing order of their respective subtree dimensions; for each sibling then, starting from the last one (LIFO), the list of its children (if any) is output repeating the same procedure recursively. E.g. one of the best topological sorts from the discourse tree of figure 3 is  $(u_1, u_4, u_5, u_2, u_3, u_6, u_8, u_7, u_9, u_{10})$ . This result is expressed in the following proposition, where algorithm 4 is a slight variation of algorithm 3.

**Proposition 3.** *Given a discourse tree with  $n$  discourse units  $u_1, u_2, \dots, u_n$  containing  $n + 1$  discourse entities  $c_0, c_1, c_2, \dots, c_n$ , the topological sorts that minimize the cost function expressed in Def. 9 first, and then the cost function of Def. 10 next, are all and only the ones returned by Algorithm 4.*

---

**Algorithm 4** Generation of the best topological sorts (Hybrid approach #1 (CT-mCD))

---

$D \leftarrow u_1$  {array  $D$  stores one by one all best linear sorts; here it is initialized with the first discourse unit}  
 $n \leftarrow$  number of nodes in tree  
 $count \leftarrow 1$  {current size of  $D$ }  
call *CalculateBestSorts*( $u_1$ )

**procedure** *CalculateBestSorts*( $u$ )

$L \leftarrow$  all children of  $u$

$L_d \leftarrow$  children of  $u$  having at least one descendant

**if**  $L_d \neq \emptyset$  **then**

$P_L \leftarrow$  list of all permutations of  $L$  where nodes in  $L_d$  come last, in decreasing order of their respective subtree dimensions;

**else**

$P_L \leftarrow$  list of all permutations of  $L$

**end if**

**for all**  $p \in P_L$  **do**

```

append array  $p$  to  $D$ 
for  $i = \text{size}(p)$  downto 1 do
    call  $\text{CalculateBestSorts}(p_i)$ 
end for
 $\text{count} \leftarrow \text{count} + \text{size}(p)$ 
if  $\text{count} = n$  then
    output  $D$ 
end if
remove array  $p$  from  $D$ 
 $\text{count} \leftarrow \text{count} - \text{size}(p)$ 
end for
end procedure

```

---

We tried then to apply in sequence the mCD-approach first, and the CT-approach next. We found out that the sequence of discourse units in each best topological sort follow the rule expressed by Proposition 4.

**Proposition 4.** *Given a discourse tree with  $n$  discourse units  $u_1, u_2, \dots, u_n$  containing  $n + 1$  discourse entities  $c_0, c_1, c_2, \dots, c_n$ , the topological sorts that minimize the cost function expressed in Def. 10 first, and the cost function of Def. 9 next, are all and only the ones that derive from a depth-first-like visit of the tree, where at any point, given an output node, we visit its subtrees ordering them by increasing size.*

The interesting feature of the outcoming orderings is that we leave long elaboration chains at the end, planning the short ones first. If we see it from the point of view of the reader, this is what she usually expects from a text describing an object: Immediate characteristics/attributes of the described object come first, and relations of this object with further entities (possibly nested) are left at the end.

We propose Algorithm 5 that calculates only one of the best orderings, since there could be more than one.

---

**Algorithm 5** Generation of one of the best topological sorts (Hybrid approach #2 (mCD-CT))

---

```

 $n \leftarrow$  number of nodes in tree
call  $\text{CalculateSort}(u_1)$ 

```

```

procedure  $\text{CalculateSort}(u)$ 
output  $u$ 
 $L \leftarrow$  all children of  $u$  sorted by increasing size of respective subtrees
for all  $p \in L$  do
    call  $\text{CalculateSort}(p)$ 
end for
end procedure

```

---

### 3.2.4 User-driven planning

If on one side the previous approaches (Centering-Theory-based, minimal conceptual distance, and hybrid) try to minimize some cost functions in order to

have a higher local coherence and/or a better distribution of entities within discourse units in the text plan, on the other side they don't take into consideration how much change is involved in the text plan whenever the user modifies the query (adding or removing branches to the query tree). The idea would be to minimize the changes in the order of the discourse units in the text plan when the user edits the query.

We could think of a text plan where the discourse units have the same order of insertion followed by the user. In this case any addition to the query is reflected in a new discourse unit appended to the text plan. E.g. if a new relation is added to concept  $c_5$  in the query tree of figure 2, let's say  $r_{10}$  along with the range concept  $c_{11}$ , this would generate the query tree of figure 5, and the discourse tree of figure 6 with the new discourse unit  $u_{11}$ .

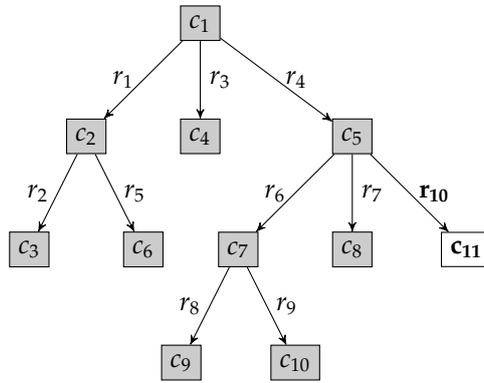


Figure 5: Adding a new relation to the query tree of Fig. 2

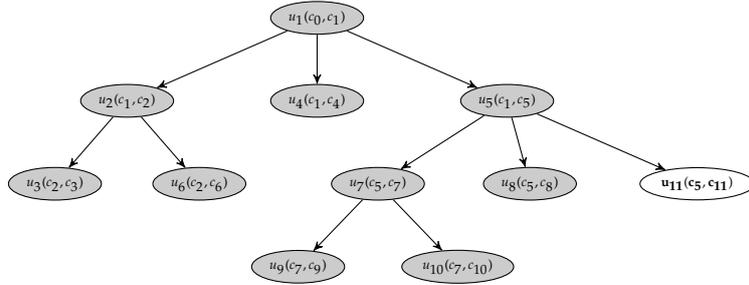


Figure 6: Discourse tree derived from the query tree of Fig. 5

The chosen topological sort, according to the planning strategy proposed, would simply be:

$$u_1(c_0, c_1), u_2(c_1, c_2), u_3(c_2, c_3), u_4(c_1, c_4), u_5(c_1, c_5), u_6(c_2, c_6), u_7(c_5, c_7), \\ u_8(c_5, c_8), u_9(c_7, c_9), u_{10}(c_7, c_{10}), u_{11}(c_5, c_{11}).$$

A newly inserted discourse unit (as  $u_{11}(c_5, c_{11})$  in the example) is always appended to the text plan, possibly far away from the latest previous discourse

unit ( $u_8(c_5, c_8)$ ), having a common discourse referent ( $c_5$ ), and farther from the unit ( $u_5(c_1, c_5)$ ) where this referent was first introduced.

Hence this kind of strategy yields on average pretty bad values in terms of local coherence and overall conceptual distance according to the measures of Definitions 9 and 10. In fact we cannot expect that the user edits the query in a coherent way, having the query tree already in mind before typing, and reproducing it immediately afterwards with a clean depth-first traversal.

The regularity of a user-driven order of discourse units, where the last inserted unit is always appended to the text plan, doesn't compensate for the bad average performance in terms of local coherence or conceptual distance.

It is evident that we need to find an appropriate trade-off between coherence/distance criteria and minimal change in the text plan after each query editing operation.

### 3.2.5 Depth-first planning

A good trade-off, slightly unbalanced in favor of the minimization of the change in the text plan after a query tree edit, is the easiest-to-obtain topological ordering, merely a *depth-first serialization* of the tree. Considering the ordering strategies discussed above, this is among the ones that can be obtained with the lowest time complexity, i.e.  $O(n)$  where  $n$  is the number of the discourse tree nodes.

Given the tree of figure 6, the topological sort according to this method would be:

$$u_1(c_0, c_1), u_2(c_1, c_2), u_3(c_2, c_3), u_6(c_2, c_6), u_4(c_1, c_4), u_5(c_1, c_5), u_7(c_5, c_7), \\ u_9(c_7, c_9), u_{10}(c_7, c_{10}), u_8(c_5, c_8), u_{11}(c_5, c_{11}).$$

Algorithm 6 shows how to obtain such simple ordering.

---

#### Algorithm 6 Generation of a topological ordering using depth-first search

---

$n \leftarrow$  number of nodes in tree  
call *CalculateSort*( $u_1$ )

**procedure** *CalculateSort*( $u$ )

output  $u$

$L \leftarrow$  all children of  $u$  from left to right

**for all**  $p \in L$  **do**

    call *CalculateSort*( $p$ )

**end for**

**end procedure**

---

To complete the list of possible planning strategies we tested, we would like to mention a further one that could be boiled down to a depth-first planning provided we change the tree topology according to further constraints. It's presented in the next subsection.

### 3.2.6 Relation-priority depth-first planning

This planning strategy, inspired by [Galanis and Androutsopoulos, 2007], requires additional information that must be pre-specified in the ontology and

attached to each relation. This information consists of ordering annotations assigned to relations, and specifies a partial order among those relations having the same domain concept. This translates into a priority value assigned to each discourse unit, and valid only locally within each set of discourse units which are siblings in the discourse tree.

The idea is to keep the query tree edges (relations) that exit from each node, sorted from left to right according to their order. In other words the sequence of relations having the same domain concept are not sorted according to the insertion order followed while creating the query, but respecting the priority value associated with each relation. It could happen though, that two or more relations have the same priority; in this case they are sorted by creation order.

The resulting discourse tree is then ready to be linearized with the simple depth-first traversal proposed above (Algorithm 6).

Figures 7 and 8 show the previous examples of query and discourse tree whose topology is modified according to the additional ordering annotations.

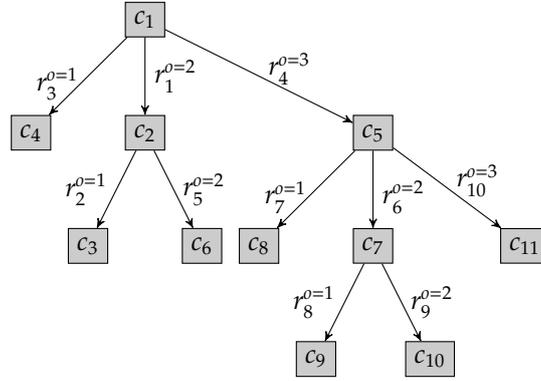


Figure 7: Query tree with ordering annotations attached to relations

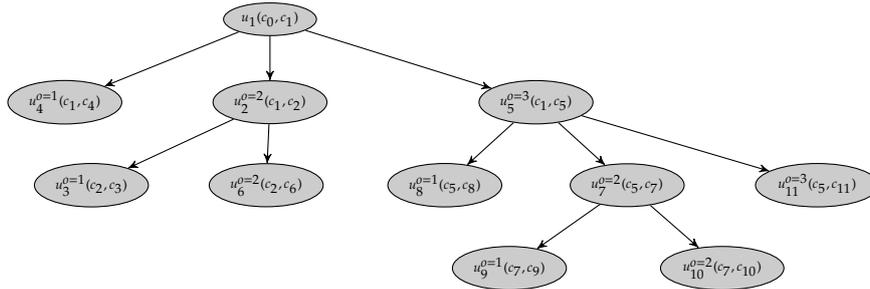


Figure 8: Discourse tree with ordering annotation attached to discourse units

Following this planning strategy, from the discourse tree of Figure 8 we obtain the following topological sort:

$$u_1(c_0, c_1), u_4(c_1, c_4), u_2(c_1, c_2), u_3(c_2, c_3), u_6(c_2, c_6), u_5(c_1, c_5), u_8(c_5, c_8), \\ u_7(c_5, c_7), u_9(c_7, c_9), u_{10}(c_7, c_{10}), u_{11}(c_5, c_{11}).$$

### 3.3 Summary

We have presented six possible strategies for discourse planning of a given complex concept description. We concentrated on three different goals

1. maximization of local referential-coherence (CT);
2. minimization of overall conceptual distance (mCD);
3. minimization of change in the discourse plan between consecutive edits (user-driven, depth-first, relation-priority depth-first).

If on one side *maximizing the referential coherence* (CT) among discourse units seemed to be a good planning strategy, on the other side we noticed that the overall conceptual proximity in the generated plans for several tree topologies was not satisfactory. We introduced then the measure of *conceptual distance* (mCD) applying it separately in a first attempt, and in hybrid approaches next. In *hybrid planning* we tried two strategies: seeking goal 1 and 2 in sequence (CT-mCD), and in reverse order (mCD-CT). This second hybrid strategy gave an interesting result, reported in Proposition 4.

While these approaches work well when we consider the complex concept description as a static input, they fail when we want that the input be created incrementally by a user requesting a sequence of plan generations. From a human-machine interaction viewpoint we would like to minimize the changes in the plan between consecutive edits. This is achieved introducing the third goal. *User-driven planning* answers this purpose but doesn't pay-off the bad average performance in terms of local coherence or conceptual distance. *Depth-first planning* although very simple, yields a better trade-off of the three goals.

We proposed then a last strategy (*relation-priority depth-first*) which requires that roles be augmented in the domain ontology with ordering annotations, specifying a partial order valid among those roles having the same domain concept. The natural rationale behind this is that when describing a concept, relations and attributes that better characterize the concept under examination should be planned first, leaving secondary roles for subsequent positions in the plan.

Assigning ordering annotations to roles can be regarded as a possible encoding of *domain communication knowledge* (DCK for short). The notion of DCK was introduced by [Kittredge *et al.*, 1991] who advanced the hypothesis that

any text planning task relies, explicitly or implicitly, on domain-specific text planning knowledge.

domain  
communication  
knowledge

The authors argue that DCK is a third kind of knowledge that natural language generators should use, at an intermediate level between *domain knowledge* and *communication knowledge*. The difference between DCK and communication knowledge is described in [Kittredge *et al.*, 1991] as follows:

Communication knowledge is independent of any particular domain knowledge. [...] Consider the task of describing a set of objects in some domain. Communication knowledge about thematic structure implies a strategy that describes together those objects that share some feature. Domain knowledge can supply information about which objects share which features. But if there are many

different features, the task remains of choosing the feature(s) according to which the descriptions will be grouped together. This choice must be based on knowledge that is neither general knowledge about communication (since the choice depends on the particular features of objects in the domain) nor actual domain knowledge (since it is only needed for planning communication).

In our context, communication knowledge is implicitly encoded in the way we traverse the query tree and therefore the deriving discourse tree. In other terms with a (relation-priority) depth first traversal we guarantee the fact that, given an output discourse plan with  $n$  discourse units, most couples of successive discourse units (maximum of  $n - 1$ ) share one discourse entity (referent).

## 4 Sentence Planning

The module that usually comes next to a text planner in most NLG systems is a *sentence planner* (otherwise called *microplanner*). It is widely recognized (even if there still is considerable debate in the NLG research community) that the main tasks of a sentence planner are

- lexicalization
- aggregation
- referring expression generation

**Lexicalization** means choosing the right words and syntactic structures to effectively communicate the message encoded in a text plan. Given the NLG system we are going to use (KPML), this process is part of the linguistic realization module, and it will be described in Section 5.

**Aggregation** deals with the quantity of information that each sentence in the text must contain.

**Referring expression generation** suggests which phrases should be used to refer to each domain entity found in the text plan.

Given these three phases, systems available to-date employ one of two possible solutions as described in [Reiter and Dale, 2000]:

- a blackboard architecture, where no specific ordering is imposed over the abovementioned phases;
- a pipelined architecture, which the various phases come in a pre-specified order.

In our system we start with sentence aggregation, followed by referring expressions generation, and finally we generate a sentence plan.

### 4.1 Sentence aggregation

Aggregation can be seen as the task of combining several input elements into a more complex structure for the sake of coherence, fluency and conciseness.

[Cheng *et al.*, 1997] give an excellent definition of aggregation that reads as follows:

Functioning as one or a set of processes acting on some intermediate text structures in text planning, aggregation decides which pieces of structures can be combined together to be realized as complex sentences later on so that a concise and cohesive text can be generated while the meaning of the text is kept almost the same as that without aggregation.

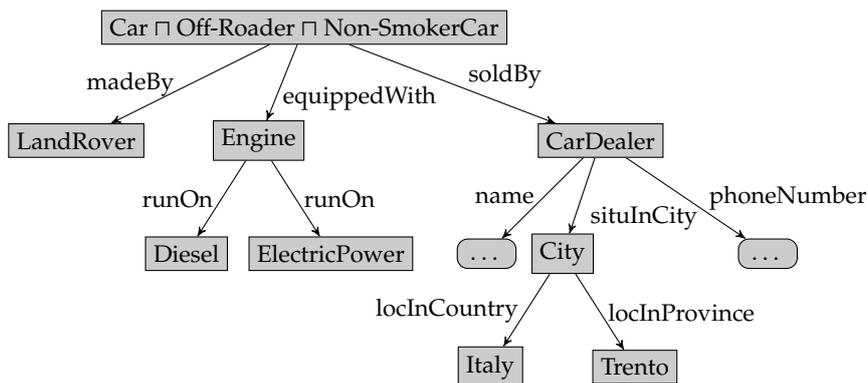


Figure 9: A query tree example.

Skipping sentence aggregation would lead to stilted texts composed by simple subject-verb-object sentences. Starting from the complex concept expressed by the graph of Figure 9, and using the relation-priority depth-first planning strategy, without aggregation we would obtain this discourse plan

$u_1, u_{11}, u_{13}, u_4, u_2, u_3, u_6, u_5, u_8, u_7, u_9, u_{10}, u_{12}$

and the following text after linguistic realization:

$u_1$  I'm looking for a car.  $u_{11}$  The car is an off-roader.  $u_{13}$  The car is a non-smoker car.  $u_4$  The car is made by Land Rover.  $u_2$  The car is equipped with an engine.  $u_3$  The engine runs on diesel.  $u_6$  The engine runs on electric power.  $u_5$  The car is sold by a car dealer.  $u_8$  The car dealer's name is [...].  $u_7$  The car dealer is situated in a city.  $u_9$  The city is in Italy.  $u_{10}$  The city is in the province of Trento.  $u_{12}$  The car dealer's phone number is [...].

The index  $i$  given to each discourse unit  $u_i$  in the example text above reflects the order of insertion followed by the user while building this query.

#### 4.1.1 Types of Aggregation

In the past twelve years, researchers working on aggregation have mainly elaborated on Wilkinson's [Wilkinson, 1995] classification which is based on *locus of process*: "Something which the various treatments of aggregation have differed on or left vague is at exactly what levels of language generation aggregation may take place. . . . In fact, aggregation-like phenomena can occur at such a variety of stages and in such a variety of ways that the term begins to seem stretched beyond its capacity." [Wilkinson, 1995]

locus of  
process

The recognized typologies of aggregation are six:

- **Conceptual aggregation:** this is the deepest locus of aggregation, where a complex concept can possibly be reduced to a simpler equivalent one by means of an inference.
- **Discourse (rhetorical) aggregation:** any operation that applies to a *discourse structure*, *rhetorical structure*, or *text plan* and maps it to a better structure or plan (how “better” must be defined by a metric).
- **Semantic aggregation:** the combination of two or more semantic entities into one by means of semantic grouping; it takes place at a level which is abstracted from syntax, but is language-dependent. [Reape and Mellish, 1999] note that they “could find no clear examples of semantic aggregation in the literature which couldn’t alternatively be classified as either conceptual, syntactic or lexical aggregation.”
- **Syntactic aggregation:** the most frequent form of aggregation. Aggregation rules that are commonly specified are (a) *subject grouping* rules and (b) *predicate grouping* rules.
- **Lexical aggregation:** includes three types of aggregation: (a) the mapping of more lexical predicates to fewer lexemes, (b) the mapping of (more) lexical predicates to (fewer) lexical predicates and (c) the mapping of (more) lexemes to (fewer) lexemes.
- **Referential aggregation:** this was introduced by [Reape and Mellish, 1999] and is not covered by [Wilkinson, 1995]; it refers to aggregation by means of referring expression generation.

In our setting, aggregation is meant to detect shared concepts and roles, and to combine them in order to reduce redundancies and repetitions in the resulting text. Given a text plan, as sequence of discourse units, we try to aggregate them according to a set of *aggregation template structures* which can be reduced to three types:

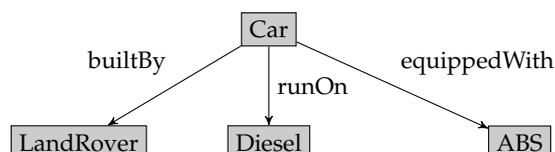
- simple conjunction structure,
- shared subject-predicate,
- syntactic embedding,

which are a subset of the aggregation roles foreseen and described in [Melenoglou, 2002] for the M-PIRO project.

The three abovementioned types mainly fall under the syntactic aggregation typology, and they are described below.

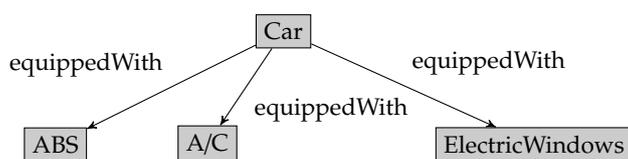
**Simple conjunction (SC)** Simple conjunction can be employed whenever we want to aggregate several roles of the same concept, and the result is the aggregation of several propositions with the same subject.

Let’s suppose we have the following relational structure:



Without aggregation the three discourse units would generate three separate sentences: *The car is built by Land Rover. The car runs on diesel. The car is equipped with ABS.* Using simple conjunction we would obtain: *The car is built by Land Rover, it runs on diesel, and it is equipped with ABS,* where for the sake of readability we pronominalized the subjects of the second and third clause (see Section 4.2 on referring expressions generation).

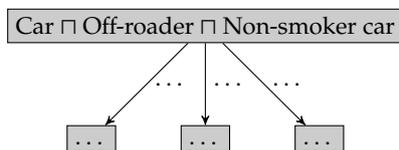
**Shared subject-predicate (SSP)** There are cases where two or more consecutive discourse units sharing the same domain concept also have the same role name. This is a case of conjunction with shared subject-predicate, as e.g. in the following relational structure:



Without aggregation we would have: *The car is equipped with ABS. The car is equipped with A/C. The car is equipped with electric windows.*

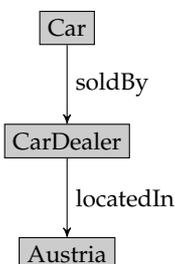
Using simple conjunction we obtain: *The car is equipped with ABS, A/C, and electric windows.*

We can also use shared subject-predicate aggregation when we need to express identity among concepts. Given a relational tree this can happen when a concept is followed by one or more compatible concepts (i.e. non-mutually-disjoint concepts) as in this example:



In aggregated form we have *The car is an off-roader and a non-smoker car.*

**Syntactic embedding (SE)** With this kind of aggregation we have a dominant proposition and a secondary proposition which is realized as a subordinated constituent as e.g. a non-defining relative clause. Starting from the following relational tree



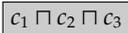
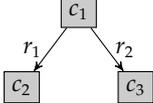
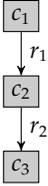
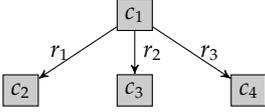
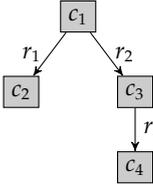
we could obtain this aggregated form: *The car is sold by a car dealer who is located in Austria.*

### 4.1.2 Aggregation Template Structures

The aggregation template structures we mentioned above are now formally listed in Table 1 sorted by number of discourse units we want to aggregate. The template structures we propose don't specify in absolute terms the maximum number of discourse units allowed ( $n$  in the generic templates below), but they impose some constraints in terms of *tree topology*, *maximum tree depth* (0–2 depending on the template), and *maximum number of different roles* (0–4 depending on the template) for each template tree.

Table 1 takes into consideration all subtree patterns we try to recognize in a given relational tree. A maximum number of aggregatable units must be defined and it represents the maximum value that can be assigned to variable  $n$ . We define with  $n_u$  the number of unique roles with the same domain concept ( $c_1$  as shown in templates n.2 and n.3).

Table 1: Aggregation template structures

Units	ID	Template	Aggregation
2	2.1		<b>shared subject-predicate</b>
	2.2		a) <b>simple conjunction</b> b) <b>shared subject-predicate</b> (if $r_1 = r_2$ )
	2.3		<b>syntactic embedding</b>
3	3.1		<b>shared subject-predicate</b>
	3.2		a) <b>simple conjunction</b> b) <b>shared subject-predicate</b> (if $r_1 = r_2 = r_3$ ) c) <b>simple conjunction</b> (between different roles) + <b>simple conjunction</b> (for roles that are equal, if either $r_1 = r_2$ or $r_2 = r_3$ )
	3.3		a) <b>simple conjunction + syntactic embedding</b> b) <b>shared subject-predicate + syntactic embedding</b> (if $r_1 = r_2$ )

*continued on next page*

Units	ID	Template	Aggregation
	3.4	<pre> graph TD   c1[c1] -- r1 --&gt; c2[c2]   c2 -- r2 --&gt; c3[c3]   c2 -- r3 --&gt; c4[c4] </pre>	<b>syntactic embedding + shared subject-predicate</b> ( $r_2 = r_3$ )
n	n.1		<b>shared subject-predicate</b>
	n.2	<pre> graph TD   c1[c1] -- r1 --&gt; c2[c2]   c1 -- r2 --&gt; c3[c3]   c1 -- ... --&gt; dots[...]   c1 -- rn --&gt; cn1[cn+1] </pre>	<ul style="list-style-type: none"> <li>a) <b>simple conjunction</b> (for <math>n \leq 3</math> and <math>r_1 \neq r_2 \neq \dots \neq r_n</math>)</li> <li>b) <b>shared subject-predicate</b> (if <math>r_1 = r_2 = \dots = r_n</math>)</li> <li>c) <b>simple conjunction</b> (between different roles, if <math>2 \leq n_u \leq 3</math>) + <b>shared subject-predicate</b> (for roles that are equal<sup>1</sup>)</li> </ul>
	n.3	<pre> graph TD   c1[c1] -- r1 --&gt; c2[c2]   c1 -- ... --&gt; dots1[...]   c1 -- rn-m-1 --&gt; cnm[cn-m]   cnm -- rn-m --&gt; cnm1[cn-m+1]   cnm -- ... --&gt; dots2[...]   cnm -- rn --&gt; cn1[cn+1] </pre>	<ul style="list-style-type: none"> <li>a) if <math>n = m + 2</math> and <math>r_2 = \dots = r_n</math>, <b>syntactic embedding + shared subject-predicate</b> (if <math>m \geq 1</math>)</li> <li>b) <b>simple conjunction</b> (if <math>n = m + 3</math> and <math>r_1 \neq r_2</math>) + <b>syntactic embedding + shared subject-predicate</b> (if <math>m \geq 1</math>)</li> <li>c) <b>shared subject-predicate</b> (if <math>n &gt; m + 2</math> and <math>r_1 = r_2 = \dots = r_{n-m-1}</math>) + <b>syntactic embedding + shared subject-predicate</b> (if <math>m \geq 1</math>)</li> <li>d) <b>simple conjunction</b> (between different roles <math>\in \{r_1, \dots, r_{n-m-1}\}</math> if <math>2 \leq n_u \leq 3</math>) + <b>shared subject-predicate</b> (for roles <math>\in \{r_1, \dots, r_{n-m-1}\}</math> that are equal) + <b>syntactic embedding + shared subject-predicate</b> (if <math>m \geq 1</math>)</li> </ul> <p>For all n.3 templates these conditions must hold: <math>m \geq 0</math> and <math>r_{n-m} = \dots = r_n</math>.</p>

For the patterns expressed above and in particular for  $n \in \{2, 3, 5, 6\}$  we show in table 2 several examples of aggregation. For  $n = 6$  and  $n = 7$  in some cases we exceed the limit of maximum number of clauses (with different roles) aggregatable in one sentence by means of simple conjunction ( $n_u = 4$ ), therefore we need to use two sentences. On the opposite, in 8.2.c even if we

<sup>1</sup>Note that the planning algorithm we have chosen (relation-priority depth-first) keeps the roles (descending from the same concept) ordered according to their priority, where multiple instances of the same role are always consecutive in the text plan, and never mixed-up with other roles having their same priority.

have 8 propositions, we are able to aggregate them all into one single sentence because  $n_u = 3$ .

Note that for the sake of readability the last column of table 2 shows the final surface form after aggregation *and* pronominalization.

Table 2: Aggregation examples

ID	Propositions	Aggregated form
2.1	The car is an off-roader. The car is a non-smoker car.	The car is an off-roader and a non-smoker car.
2.2.a	The car is a Land Rover. It is equipped with ABS.	The car is a Land Rover and it's equipped with ABS.
2.2.b	The car is equipped with air conditioning. The car is equipped with electric windows.	The car is equipped with air conditioning and electric windows.
2.3.a	The car is equipped with an engine. The engine runs on diesel.	The car is equipped with an engine that runs on diesel.
3.1	The car is an off-roader. The car is a demonstration car. The car is a non-smoker car.	The car is an off-roader, a demonstration car, and a non-smoker car.
3.2.a	The car is a Land Rover. Its model is Defender. It is equipped with a traction control system.	The car is a Land Rover, its model is Defender, and it is equipped with a traction control system.
3.2.b	The car is equipped with ABS. The car is equipped with air conditioning. The car is equipped with electric windows.	The car is equipped with ABS, air conditioning, and electric windows.
3.2.c	The car is a Land Rover. The car is equipped with ABS. The car is equipped with air conditioning.	The car is a Land Rover and it's equipped with ABS and air conditioning.
3.3.a	The car is an off-roader. It is equipped with an engine. The engine runs on diesel.	The car is an off-roader and it's equipped with an engine that runs on diesel.
3.3.b	It is equipped with ABS. The car is equipped with an engine. The engine runs on diesel.	The car is equipped with ABS and an engine that runs on diesel.
3.4	The car is equipped with an engine. The engine runs on gasoline. The engine runs on methane.	The car is equipped with an engine that runs on gasoline and methane.
5.1	<i>similar to 3.1</i>	
5.2.a	<i>similar to 3.2.a</i>	
5.2.b	<i>similar to 3.2.b</i>	
5.2.c	The car is a Land Rover. The car's model is Defender. The car is equipped with ABS. The car is equipped with air conditioning. The car is equipped with electric windows.	The car is a Land Rover, its model is Defender and it's equipped with ABS and air conditioning.

*continued on next page*

ID	Propositions	Aggregated form
5.3.b	The car is a Land Rover. The car's model is Defender. The car is equipped with ABS. The car is equipped with air conditioning. The car is equipped with an engine. The engine runs on diesel.	The car is a Land Rover, its model is Defender and it's equipped with ABS and an engine that runs on diesel.
5.3.c	<i>similar to 3.3.b</i>	
5.3.d	The car is a Land Rover. The car's model is Defender. The car is equipped with an engine. The engine runs on gasoline. The engine runs on methane.	The car is a Land Rover, its model is Defender, and it's equipped with an engine that runs on gasoline and methane.
6.2.d $n_u = 4$	The car is a Land Rover. The car's model is Defender. The car's color is yellow. The car is equipped with ABS. The car is equipped with air conditioning. The car is equipped with electric windows.	The car is a Land Rover, its model is Defender and it's color is yellow. It's equipped with ABS, air conditioning, and electric windows.
7.2.c $n_u = 4$	The car is a Land Rover. The car's model is Defender. The car's color is yellow. The car's color is blue. The car is equipped with ABS. The car is equipped with air conditioning. The car is equipped with electric windows.	The car is a Land Rover, its model is Defender and its color is yellow and blue. It's equipped with ABS, air conditioning, and electric windows.
8.2.c	The car is a Land Rover. The car's color is yellow. The car's color is blue. The car's color is red. The car is equipped with a traction control system. The car is equipped with ABS. The car is equipped with air conditioning. The car is equipped with electric windows.	The car is a Land Rover, its color is yellow, blue, and red, and it's equipped with a traction control system, ABS, and air conditioning.

The next step is meant to find a way to compute the best match of the patterns of table 1 for a given relational tree taking into consideration that the text planning algorithm we choose is the *relation-priority depth-first*.

#### 4.1.3 Best template structure matching

The idea underlying this is to linearize both the relational tree (according to the chosen planning algorithm) and the template structures, seeking the best covering match of the linearized templates in the plan that minimizes the number of sentences in the outcoming sentence plan.

The linearized templates we foresee are listed in table 3. Concepts are represented as  $C_i, C_{i+1}, C_{i+2}, \dots$  where the emphasized index ( $i, i+1, i+2, \dots$ ) stands for the level in the tree where the concept (as node) is situated. Roles are represented as  $R_j, R_{j+1}, R_{j+2}, \dots$  where the emphasized index ( $j, j+1, j+2, \dots$ ) is the same for edges of the tree that represent the same role.

Table 3 groups the linearization of the allowed tree templates ordered by increasing number of constituents; it also shows that each linearization corresponds to one or more tree templates.

#	Linearized template	Original ID
1	$C_i,$	n.1
2	$(C_i R_j C_{i+1},)_+$	n.2.b
3	$(C_i R_j C_{i+1},) + (C_{i+1} R_{j+1} C_{i+2},)_+$	n.3.a, n.3.b, n.3.c
4	$(C_i R_j C_{i+1},) + (C_i R_{j+1} C_{i+1},)_+$	n.2.c
5	$(C_i R_j C_{i+1},) + (C_i R_{j+1} C_{i+1},) + (C_{i+1} R_{j+2} C_{i+2},)_+$	n.3.d
6	$(C_i R_j C_{i+1},) + (C_i R_{j+1} C_{i+1},) + (C_i R_{j+2} C_{i+1},)_+$	n.2.a, n.2.c
7	$(C_i R_j C_{i+1},) + (C_i R_{j+1} C_{i+1},) + (C_i R_{j+2} C_{i+1},) + (C_{i+1} R_{j+3} C_{i+2},)_+$	n.3.d

Table 3: Linearized templates

In order to find the best template structure match, we need to convert the text plan output by the previous phase using the same notation employed for the linearizations above.

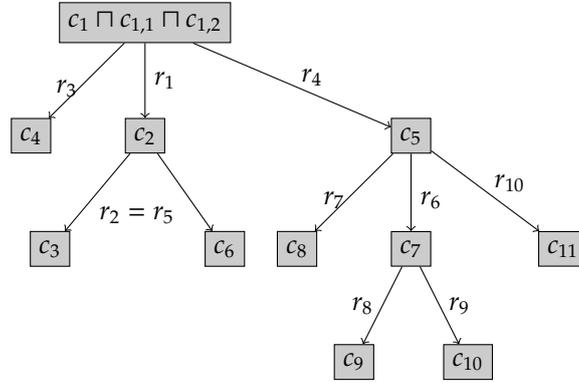


Figure 10: A query tree waiting to be linearized

Let's suppose we start from the tree of Figure 9 remapped to the one in Figure 10 for the sake of simplicity. Provided we use the planning algorithm proposed in section 3.2.6 (relation-priority depth-first planning), the text plan with the notation introduced above would be:

$$C1C1C1, C1R1C2, C1R2C2, C2R3C3, C2R3C3, C1R4C2, C2R5C3, C2R6C3, \\ C3R7C4, C3R8C4, C2R9C3, \quad (1)$$

The list of templates needs to be instantiated: indexes  $i$  and  $j$  must be initialized according to the index (level) of the first concept and the first role in the text plan respectively. We obtain a list of regular expressions we call *aggregation patterns* or simply *patterns*.

1	(C1)+,
2	(C1R1C2,)+
3	(C1R1C2,)+(C2R2C3,)+
4	(C1R1C2,)+(C1R2C2,)+
5	(C1R1C2,)+(C1R2C2,)+(C2R3C3,)+
6	(C1R1C2,)+(C1R2C2,)+(C1R3C2,)+
7	(C1R1C2,)+(C1R2C2,)+(C1R3C2,)+(C2R4C3,)+

The pattern matching is done starting from the longest pattern #7 to the shortest one #1. The first pattern that matches is #1 followed by #5. We are able to aggregate the first three compatible concepts (C1C1C1,) in a first sentence, followed by another sentence that aggregates four units (C1R1C2, C1R2C2, C2R3C3, C2R3C3,). After these first two hits, no other pattern matches the remaining part of the plan:

C1R4C2, C2R5C3, C2R6C3, C3R7C4, C3R8C4, C2R9C3,

We need to re-instantiate the list of templates setting  $i = 1$  and  $j = 4$ . The patterns become then:

1	(C1+),
2	(C1R4C2,)+
3	(C1R4C2,)+(C2R5C3,)+
4	(C1R4C2,)+(C1R5C2,)+
5	(C1R4C2,)+(C1R5C2,)+(C2R6C3,)+
6	(C1R4C2,)+(C1R5C2,)+(C1R6C2,)+
7	(C1R4C2,)+(C1R5C2,)+(C1R6C2,)+(C2R8C3,)+

This time there is only one pattern that matches, namely #3. The part of the text plan that remains to be matched is

C2R6C3, C3R7C4, C3R8C4, C2R9C3,

The templates need to be instantiated again with  $i = 2$  and  $j = 6$  yielding these patterns:

1	(C2)+,
2	(C2R6C3,)+
3	(C2R6C3,)+(C3R7C4,)+
4	(C2R6C3,)+(C2R7C3,)+
5	(C2R6C3,)+(C2R7C3,)+(C3R8C4,)+
6	(C2R6C3,)+(C2R7C3,)+(C2R8C3,)+
7	(C2R6C3,)+(C2R7C3,)+(C2R8C3,)+(C3R10C4,)+

Again, pattern #3 is the only one that matches. There are two more units to be matched:

C3R8C4, C2R9C3,

This time we need to instantiate just the first four templates (whose minimum length doesn't exceed the remaining two units), with  $i = 3$  and  $j = 8$ :

1	C3(C3)+,
2	(C3R8C4,)+
3	(C3R8C4,)+(C4R9C5,)+
4	(C3R8C4,)+(C3R9C4,)+

where only pattern #2 matches. Finally the last set of patterns is generated setting  $i = 2$  and  $j = 9$ , in order to match the very last unit with pattern #2 shown below.

C2R9C3,

1	C2(C2)+,
2	(C2R9C3,)+

Summarizing, given the query tree of Figure 10, the text plan resulting from it is composed by 12 clauses which, according to the proposed templates, can be joined to form 6 sentences ( $S_1 \dots S_6$ ) as reported below:

Sentences	Patterns	
$S_1$ $c_1 r_{is} c_{1,1} \oplus_{ssp} c_1 r_{is} c_{1,2}$	C1C1C1,	#1
$S_2$ $c_1 r_3 c_4 \oplus_{sc} c_1 r_1 c_2 \oplus_{se} c_2 r_2 c_3 \oplus_{ssp} c_2 r_5 c_6$	C1R1C2, C1R2C2, C2R3C3, C2R3C3,	#5
$S_3$ $c_1 r_4 c_5 \oplus_{se} c_5 r_7 c_8$	C1R4C2, C2R5C3,	#3
$S_4$ $c_5 r_6 c_7 \oplus_{se} c_7 r_8 c_9$	C2R6C3, C3R7C4,	#3
$S_5$ $c_7 r_9 c_{10}$	C3R8C4,	#2
$S_6$ $c_5 r_{10} c_{11}$	C2R9C3,	#2

At this point a few comments on the notation used are necessary. The three kinds of aggregation we use, namely *simple conjunction*, *shared subject-predicate*, and *syntactic embedding* are represented with the three symbolic operators  $\oplus_{sc}$ ,  $\oplus_{ssp}$ , and  $\oplus_{se}$ . The conjunction of concepts  $c_{1,1}$  and  $c_{1,2}$  with  $c_1$  has been splitted into two clauses aggregated by means of a shared subject-predicate ( $\oplus_{ssp}$ ). The newly introduced *compatibility relation*  $r_{is}$  accounts for these conjunctions i.e. that  $c_{1,1}$  and  $c_{1,2}$  must be compatible (not disjoint) with  $c_1$ . This relation will be lexicalized with a copula as shown in the examples *n.1* of Table 2.

$r_{is}$   
compatibility  
relation

Using the original concept and role names of Figure 9, the resulting 6 sentences would be:

- $S_1$ : Car·is·Off-roader  $\oplus_{ssp}$  Car·is·Non-SmokerCar
- $S_2$ : Car·madeBy·Land Rover  $\oplus_{sc}$  Car·equippedWith·Engine  $\oplus_{se}$  Engine·runOn·Diesel  $\oplus_{ssp}$  Engine·runOn·ElectricPower
- $S_3$ : Car·soldBy·CarDealer  $\oplus_{se}$  CarDealer·name·[...]
- $S_4$ : CarDealer·situatedInCity·City  $\oplus_{se}$  City·locInCountry·Italy
- $S_5$ : City·locInProvince·Trento
- $S_6$ : CarDealer·phoneNumber·[...]

#### 4.1.4 Algorithms

In this section we describe two algorithms: The first one (Algorithm 7) is needed to obtain the text plan with the notation introduced in the previous section; the second one (Algorithm 8) finds the best covering match of the text plan using the aggregation templates of table 3.

---

**Algorithm 7** Linearization of a query tree

---

```
 $u_1 \leftarrow$  root node  
 $p \leftarrow$  "" {text plan as string}  
 $P \leftarrow$  empty vector {text plan as vector of pointers to query entities}  
 $i_r \leftarrow 1$  {role counter}  
 $r_{-1} \leftarrow$  null {previous role}  
call CalculateTextPlan( $u_1$ )
```

```
procedure CalculateTextPlan( $u$ )  
append "C" & level( $u$ ) to  $p$   
append main concept of  $u$  to  $P$ ;  
 $C_u \leftarrow$  list of compatible concepts in node  $u$   
for all  $c \in C_u$  do  
    append "C" & level( $u$ ) to  $p$   
    append  $c$  to  $P$   
end for  
append ", " to  $p$   
 $R_u \leftarrow$  all edges (roles) from  $u$  {left to right}  
for all  $r \in R_u$  do  
     $v \leftarrow$  target node of edge  $r$   
    {same roles ( $r = r_{-1}$ ) keep the same index  $i_r$  in the new notation}  
    if  $r \neq r_{-1}$  then  
         $i_r \leftarrow i_r + 1$   
    end if  
    append "C" & level( $u$ ) & "R" &  $i_r$  & "C" & level( $u$ ) + 1 & ", " to  $p$   
    append main concept of  $u$  to  $P$   
    append  $r$  to  $P$   
    append main concept of  $v$  to  $P$   
     $r_{-1} \leftarrow r$   
    call CalculateTextPlan( $v$ )  
end for  
end procedure
```

---

The outputs of this algorithm both represent the text plan as serialization of a given query tree:

- a string  $p$  with the notation shown above;
- a vector  $P$  of pointers to the entities composing the query.

Algorithm 8, instead, takes the output string  $p$  and calculates the best covering match of the text plan using the aggregation templates of table 3. The templates are instantiated at the beginning according to the indexes assigned to the first concept and the first role of  $p$ . The resulting patterns are matched

against  $p$  starting from the longest to the shortest one. The matching template number and the match are saved; then the match is removed from  $p$ . As soon as none of the patterns matches, the templates are instantiated again as above. The process stops when  $p$  is empty. The output consists of two lists: a list containing the sequence of matches, and a list of the template IDs corresponding to each match.

---

**Algorithm 8** Calculation of the best covering match

---

```

input  $T$  {list of templates to be instantiated}
input  $p$  {text plan as string}
input  $P$  {text plan as vector of pointers to query entities}
 $p_{tmp} \leftarrow p$  {tmp copy of text plan}
 $T_{inst} \leftarrow$  empty list {list of patterns as instances of  $T$ }
 $L_m \leftarrow$  empty list {list of string matches}
 $L_t \leftarrow$  empty list {list of template IDs that matched}
 $M \leftarrow$  empty list {array containing in  $M[0]$  the match, in  $M[1]$  the ID of the
matching pattern}
 $i \leftarrow$  index (level) of first concept in  $p_{tmp}$ 
 $j \leftarrow$  index of first role in  $p_{tmp}$ 
 $T_{inst} \leftarrow$  generatePatterns( $T, i, j$ )
while  $p_{tmp}$  not empty do
   $M \leftarrow$  getMatch( $T_{inst}, \&p_{tmp}$ )
  if  $M = null$  then
     $i \leftarrow$  index (level) of first concept in  $p_{tmp}$ 
     $j \leftarrow$  index of first role in  $p_{tmp}$ 
     $T_{inst} \leftarrow$  generatePatterns( $T, i, j$ )
     $M \leftarrow$  getMatch( $T_{inst}, \&p_{tmp}$ )
  end if
  append  $M[0]$  to  $L_m$  {match}
  append  $M[1]$  to  $L_t$  {template ID}
end while
output  $L_m$ 
output  $L_t$ 

function generatePatterns( $T, i, j$ )
for all  $t \in T$  do
  substitute  $i, j$  in  $t$ 
  calculate indexes of concepts C and roles R in  $T$ 
  append  $t$  to  $T_{inst}$ 
end for
return  $T_{inst}$ 
end function

function getMatch( $T_{inst}, p_{tmp}$ )
 $M \leftarrow null$ 
for  $i = \text{length}(T_{inst})$  downto 1 do
  if existsMatch( $T_{inst}[i], p_{tmp}$ ) then
     $M[0] \leftarrow$  match  $T_{inst}[i]$  in  $p_{tmp}$ 
     $M[1] \leftarrow i$ 
    remove  $M[0]$  from  $p_{tmp}$ 

```

```
    break
  end if
end for
return M
end function
```

---

## 4.2 Referring expressions generation

Referring expressions represent the ways we can consider to refer to discourse entities in a message or text in general. As [Reiter and Dale, 2000] clearly explain, *the symbolic names of knowledge base entities within these messages need to be replaced by the semantic content of noun phrase referring expressions that will be sufficient to identify the intended referents to the hearer*. The reference to a discourse entity can be done by means of a noun phrase in several ways:

1. **definite noun phrases** (as e.g. *'the car'*): these are used when referring to an entity that has already been introduced before, or when the entity is assumed to be known or inferable by the hearer;
2. **indefinite noun phrases** (as e.g. *'a car'*): this is the case when we refer to a new discourse entity that hasn't been previously mentioned;
3. **definite pronouns** (he, she, it, ...) usually anaphoric<sup>2</sup>, and typically referring to entities mentioned in the same or the previous sentence;
4. **indefinite pronouns** (one, as in *'the regular one'*);
5. **relative pronouns as subject** (who, that, which), referring to an entity contained in the previous clause;
6. **names**, where named entities can be referred to using portions of their name (*'The writer Richard Wright' → 'Richard Wright'*)

Of the above categories, we restrict the generation of referring expressions to definite noun phrases (as subject), indefinite noun phrases (as direct or indirect object), definite pronouns (as subject), and relative pronouns (as subject).

We also report the use of

- **possessive pronouns** when referring to one of the attributes of a previously mentioned entity (e.g. *The engine's displacement size is 2500 cc, and its weight is 250 kg*);
- **relative pronouns used as possessives** (like *whose*), to incorporate a reference to the possessor of an attribute following the pronoun. The possessor is usually introduced in the previous clause within the same sentence (e.g. *I'm looking for a car whose make is Lada*).

---

<sup>2</sup>A reference is said to be *anaphoric* if its interpretation depends on a preceding entity in the discourse, which is called the *antecedent*.

We start by listing some constraints we have to take into account during this phase.

The first and most general constraint is that all entities of the text plan (except the subject of the first unit which is in first person singular form) will be rendered in third person singular form.

Moreover, for each one of the referring expressions we use, there are certain constraints we have to respect that limit the position that the expression can occupy within a sentence:

**indefinite noun phrases (R-INP)** are always in (direct or indirect) object position, and they are used the first time an entity appears in the text;

**definite noun phrases (R-DNP)** are always in subject position; otherwise this would mean that the entity, being also in object position the first time it was mentioned in the text, is co-referenced by two roles, which is impossible for our definition of conjunctive query (see Section . . .).

**definite pronouns (R-DP)** are always in subject position; in this case we must be careful to respect the gender of the referent;

**relative pronouns as subject (R-RPS)** which must be the same as the object of the previous unit;

**possessive pronouns (R-PP)** can only precede a subject; they must refer to the subject of the previous unit, not to the object, otherwise

**relative pronouns as possessives (R-RPP)** would be the right choice.

Given these constraints it turns out to be very easy to assign the first two referring expressions: for each discourse unit, the first entity is tagged as a *definite noun phrase* (R-DNP) and the second as an *indefinite noun phrase* (R-INP). At this point we have to note that this pre-assignment of a definite or indefinite status to entities will not affect those entities that will be lexicalized either as proper nouns or uncountable nouns. We will see this further on, when we handle the generation of sentence plans.

From this point on, the task is to deal with the pronominalization of the first entity of each unit. We could easily borrow the idea of the local focus of attention, in particular the pronominalization strategy proposed by Centering Theory [Grosz *et al.*, 1995], which states in **Rule 1** that

If any element of  $C_f(U_n)$  is realized by a pronoun in  $U_{n-1}$ , then the  $C_b(U_{n+1})$  must be realized by a pronoun also.

In other terms, citing again the authors, this means that

[. . .] no element in an utterance can be realized as a pronoun unless the backward-looking center of the utterance is realized as a pronoun also.

where utterance ( $U_n$ ) is what we call discourse unit or simply unit (with a lower-case notation  $u_n$ ).

This rule, though, does not discern among the four categories of pronouns we have, indicating which one we should use. In principle we could simply

use *definite pronouns* (R-DP), but we want to go beyond the simple achievement of grammatical sentences, having a higher degree of fluency, conciseness, and avoiding repetitions.

The previous phase (see Section 4.1) yielded the aggregation (where possible) of several discourse units into what will become multi-clausal sentences. Within the same sentence we can have clauses (units) whose first entity (subject) is the same as the second entity (object) of the previous unit. This is a case in which the pronoun of the latter unit is a *relative pronoun as subject* (R-RPS), of what will become a relative clause.

If the role expressed in a unit is concrete (i.e. an attribute of the first entity), and the first entity of the current unit is the same as the first entity of the previous unit, the role will be the subject, prepended by a *possessive pronoun* (R-PP). The two consecutive units don't need to be part of the same sentence.

Finally, in the same setting as the previous paragraph, where instead the first entity of the current unit has to correspond to the second entity of the previous unit, we would prepend the role (subject) of the present unit with a *relative pronoun used as possessive* (R-RPP).

Table 4 reports an example for each one of the referring expressions we took into consideration.

Table 4: Examples of usage of referring expressions

Ref. expr.	Query tree	Sentence
R-INP	<pre> graph TD   I[I] -- lookFor --&gt; Car[Car]           </pre>	I'm looking for <u>a car</u> .
R-DNP	<pre> graph TD   Car[Car] -- equippedWith --&gt; ABS[ABS]           </pre>	<u>The car</u> is equipped with ABS.
R-DP	<pre> graph TD   Car[Car] -- madeBy --&gt; FIAT[FIAT]   Car -- equippedWith --&gt; ABS[ABS]           </pre>	The car is made by FIAT and <u>it</u> is equipped with ABS.

*continued on next page*

Ref. expr.	Query tree	Sentence
R-RPS	<pre> graph TD   Car[Car] -- soldBy --&gt; CarDealer[CarDealer]   CarDealer -- locatedIn --&gt; Italy[Italy] </pre>	The car is sold by a car dealer <u>who</u> is located in Italy.
R-PP	<pre> graph TD   Car[Car] -- model --&gt; Bravo[Bravo 1.6]   Car -- color --&gt; blue[blue] </pre>	The car's model is Bravo 1.6 and <u>its</u> color is blue.
R-RPP	<pre> graph TD   Car[Car] -- equippedWith --&gt; Engine[Engine]   Engine -- power --&gt; Dots[...] </pre>	The car is equipped with an engine <u>whose</u> power is . . . .

We condense now all the previous considerations, rules, and constraints into an algorithm for the generation of referring expressions. The input of the algorithm is an ordered list of all entities we have in our text plan, with the additional aggregation information obtained from the aggregation algorithm.

The output will be the same list of entities, where each entity will be completed with additional information about the referring expression to be used.

To accomplish this task we need a few functions:

- $\text{getUnit}(c_i)$  returns the discourse unit where entity  $c_i$  is to be found;
- $\text{getPreviousUnit}(u_k)$  returns the unit preceding  $u_k$ ;
- $\text{getPreviousEntity}(c_i)$  returns the entity preceding  $c_i$ ;
- $\text{getFirstEntity}(u_k)$  returns the first entity in  $u_k$ ;
- $\text{getNextEntity}(c_i)$  returns the next (to the current one) entity in  $u_k$ ;
- $\text{getLastEntity}(u_k)$  returns the last entity  $u_k$ ;
- $\text{getPosition}(c_i)$  returns the relative position of entity  $c_k$  within its discourse unit; the position is a positive integer in  $\{1, 2\}$  for units like  $c_j r_k c_l$  or an integer in a bigger set  $\{1, 2, 3, \dots\}$  for units such as  $c_i \sqcap c_{i,1} \sqcap c_{i,2} \sqcap \dots$  which represent the conjunction of two or more compatible concepts;

- `getSentence( $u_k$ )` returns the sentence to which unit  $u_k$  has been assigned after aggregation;
- `inSameSentence( $u_k, u_l$ )` which returns true if the two discourse units  $u_k$  and  $u_l$  are part of the same sentence after aggregation, otherwise it returns false;
- `sameConcept( $c_i, c_j$ )` returns true if the two entities refer to the same concept;
- `setRefExpr( $c_i, refExpr$ )` sets the given referring expression `refExpr` in  $c_i$ ;
- `existsRole( $c_i, c_j$ )` returns true if  $c_i$  and  $c_j$  are connected by a role, false otherwise;
- `isConcreteRole( $c_i, c_j$ )` returns true if the role having  $c_i$  as domain and  $c_j$  as range is a *concrete role* (attribute), false if it is an *abstract role* (relation).

---

**Algorithm 9** Generation of appropriate referring expressions for each entity present in a given text plan

```

input  $P$  {text plan as vector of discourse entities, which are uniquely identified, even though it can happen that two entities refer to the same KB concept}
for all  $c \in P$  do
   $u_{cur} \leftarrow getUnit(c)$ 
   $u_{prev} \leftarrow getPreviousUnit(u_{cur});$ 
  if  $getEntityPosition(c) = 1$  then
     $c_{next} = getNextEntity(c)$ 
     $setRefExpr(c, R-DNP)$ 
    if  $u_{prev} \neq NULL$  then
      if  $sameConcept(getFirstEntity(u_{prev}), c)$  then
         $setRefExpr(c, R-DP)$ 
        if  $existsRole(c, c_{next})$  then
          if  $isConcreteRole(c, c_{next})$  then
             $setRefExpr(c, R-PP)$ 
          end if
        end if
      else if  $inSameSentence(u_{prev}, u_{cur})$  then
        if  $sameConcept(getLastEntity((u_{prev}), c))$  then
           $setRefExpr(c, R-RPS)$ 
          if  $existsRole(c, c_{next})$  then
            if  $isConcreteRole(c, c_{next})$  then
               $setRefExpr(c, R-RPP)$ 
            end if
          end if
        end if
      end if
    else
       $setRefExpr(c, R-INP)$ 
    end if
  end for

```

---

A few comments at this point are necessary. First of all, the use of possessive pronouns (R-PP) and relative pronouns as possessives (R-RPP) is not restricted to entities connected with a concrete role: We can have cases where an entity is followed by an abstract role which behaves as an attribute and is therefore rendered as a substantive instead of a predicate, as abstract roles usually are.

For example, the query of figure 11 would be rendered as “I am looking for a car whose make is Santana”, where the abstract role *make* is rendered as the subject of the second unit, and the reference to car is incorporated into the relative pronoun (*whose*).

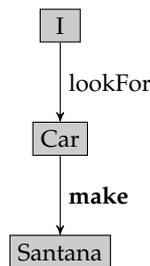


Figure 11: Query with abstract role (*make*) that is rendered as substantive.

This is to be considered an exception, since the rule is that abstract roles are usually rendered as a predicate (as e.g. the abstract role *lookFor*).

Another issue regards the correct choice of a pronoun according to the gender of the *referent* (third person singular), and the fact that *they* are either human or non human entities. The problem arises when we want to refer to a single definite person androgynously, i.e. with a gender-neutral pronoun. There are various viable solutions. We could try to **avoid using the pronoun**, but this would lead to annoying repetitions of the name that should have been pronominalized. In order to avoid sexist writing we could **alternate male and female pronouns**: in this case this would be pretty confusing for the user. We very often see people using **both pronouns together** but this is considered by readers and writers stylistically inelegant. Excluding the possibility of **inventing a new pronoun**, what remains—and this is the solution we adopt—is resorting to plural pronouns such as *they*, and *their* for singular uses. This is called the **singular they**.

singular they

*Singular they* is a popular, non-technical expression for uses of the pronoun *they* (and its inflected forms) when plurality is not required by the context. Singular *they* remains morphologically and syntactically plural, and its use as pronoun of indefinite gender and indefinite number is well established in speech and writing, even in literary and formal contexts [Merriam-Webster, 2007]. We weaved an example of *singular they* usage in the previous paragraph, in correspondence of the margin note.

The assignment of the correct pronoun to each pronominalizable entity will be dealt in detail in the following section. We only anticipate in Table 5 the set of all pronouns we are going to use.

We recap now the query example of Figure 9 and the successive aggregation result on page 33. The aggregation output is completed with the first

Type	Non-human	Human		
	Indefinite	Masculine	Feminine	Indefinite
<b>R-DP</b>	it	he	she	they
<b>R-RPS</b>	that	who/that		
<b>R-PP</b>	its	his	her	their
<b>R-RPP</b>	whose			

Table 5: Complete set of singular pronouns used

introductory sentence  $S_0$  and the discourse unit labels  $u_1, u_2, \dots, u_{13}$ .

$S_0: u_1(\text{I}\cdot\text{lookFor}\cdot\text{Car})$

$S_1: u_{11}(\text{Car}\cdot\text{is}\cdot\text{Off-roader}) \oplus_{\text{ssp}} u_{13}(\text{Car}\cdot\text{is}\cdot\text{Non-SmokerCar})$

$S_2: u_4(\text{Car}\cdot\text{madeBy}\cdot\text{Land Rover}) \oplus_{\text{sc}} u_2(\text{Car}\cdot\text{equippedWith}\cdot\text{Engine}) \oplus_{\text{se}} u_3(\text{Engine}\cdot\text{runOn}\cdot\text{Diesel})$   
 $\oplus_{\text{ssp}} u_6(\text{Engine}\cdot\text{runOn}\cdot\text{ElectricPower})$

$S_3: u_5(\text{Car}\cdot\text{soldBy}\cdot\text{CarDealer}) \oplus_{\text{se}} u_8(\text{CarDealer}\cdot\text{name}\cdot[\dots])$

$S_4: u_7(\text{CarDealer}\cdot\text{situatedInCity}\cdot\text{City}) \oplus_{\text{se}} u_9(\text{City}\cdot\text{locInCountry}\cdot\text{Italy})$

$S_5: u_{10}(\text{City}\cdot\text{locInProvince}\cdot\text{Trento})$

$S_6: u_{12}(\text{CarDealer}\cdot\text{phoneNumber}\cdot[\dots])$

After running Algorithm 9 on all discourse units except the first one ( $u_1$ ), each discourse entity is assigned an appropriate referring expression label. The output is:

$S_0: u_1(\text{I}\cdot\text{lookFor}\cdot\text{Car})$

$S_1: u_{11}(\text{Car}_{\text{R-DP}}\cdot\text{is}\cdot\text{Off-roader}_{\text{R-INP}}) \oplus_{\text{ssp}} u_{13}(\text{Car}_{\text{R-DP}}\cdot\text{is}\cdot\text{Non-SmokerCar}_{\text{R-INP}})$

$S_2: u_4(\text{Car}_{\text{R-DP}}\cdot\text{madeBy}\cdot\text{Land Rover}_{\text{R-INP}}) \oplus_{\text{sc}} u_2(\text{Car}_{\text{R-DP}}\cdot\text{equippedWith}\cdot\text{Engine}_{\text{R-INP}})$   
 $\oplus_{\text{se}} u_3(\text{Engine}_{\text{R-RPS}}\cdot\text{runOn}\cdot\text{Diesel}_{\text{R-INP}}) \oplus_{\text{ssp}} u_6(\text{Engine}_{\text{R-DP}}\cdot\text{runOn}\cdot\text{ElectricPower}_{\text{R-INP}})$

$S_3: u_5(\text{Car}_{\text{R-DNP}}\cdot\text{soldBy}\cdot\text{CarDealer}_{\text{R-INP}}) \oplus_{\text{se}} u_8(\text{CarDealer}_{\text{R-RPP}}\cdot\text{name}\cdot[\dots]_{\text{R-INP}})$

$S_4: u_7(\text{CarDealer}_{\text{R-DP}}\cdot\text{situatedInCity}\cdot\text{City}_{\text{R-INP}}) \oplus_{\text{se}} u_9(\text{City}_{\text{R-RPS}}\cdot\text{locInCountry}\cdot\text{Italy}_{\text{R-INP}})$

$S_5: u_{10}(\text{City}_{\text{R-DP}}\cdot\text{locInProvince}\cdot\text{Trento}_{\text{R-INP}})$

$S_6: u_{12}(\text{CarDealer}_{\text{R-DNP}}\cdot\text{phoneNumber}\cdot[\dots]_{\text{R-INP}})$

### 4.3 Generation of a Sentence Plan in SPL

With the outputs obtained from the discourse planning, sentence aggregation, and referring expression generation phases, we are ready to generate the input for the linguistic realizer. The input is called *sentence plan* and the language used is the *Sentence Plan Language* or simply SPL, a language devised by Robert

Kasper [Kasper, 1989]. The details of this formalism are thoroughly explained in Section 6.4.

In short, SPL is the form of non-linguistic input adopted by several linguistic realizers, among which we mention  $\kappa\text{PML}$  [Bateman, 1997a], the one we adopted. In a more general way we can say that an SPL is the semantic specification of a sentence.

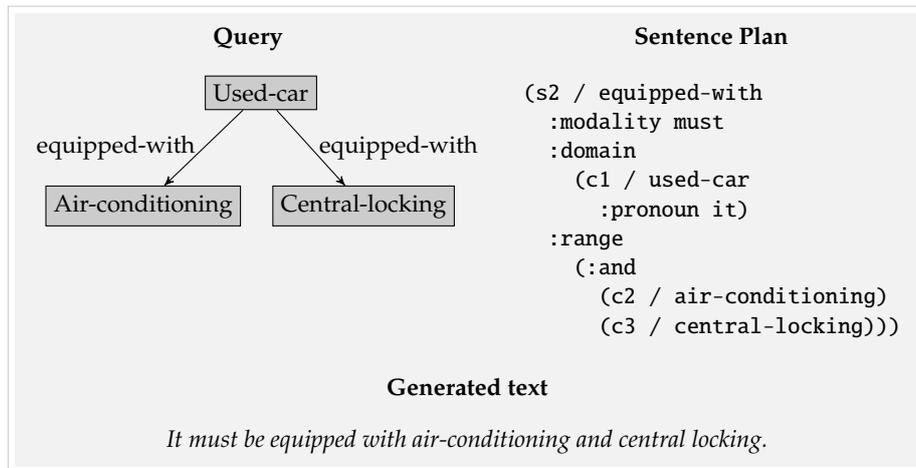
We start with some examples, where for each one we show the query, the sentence plan it is mapped into, and the  $\kappa\text{PML}$ -generated text.

**Example 1** In this first example the query is a conjunction of three compatible concepts: *Used-car*, *Off-Roader*, and *Non-smoker-car*.

Query	Sentence Plan
Used-car $\sqcap$ Off-roader $\sqcap$ Non-smoker-car	(s1 / class-ascription :modality must :domain (c1 / used-car :determiner the) :range (:and (c2 / off-roader) (c3 / non-smoker-car)) )
	<b>Generated text</b>
	<i>The used-car must be an off-roader and a non-smoker car.</i>

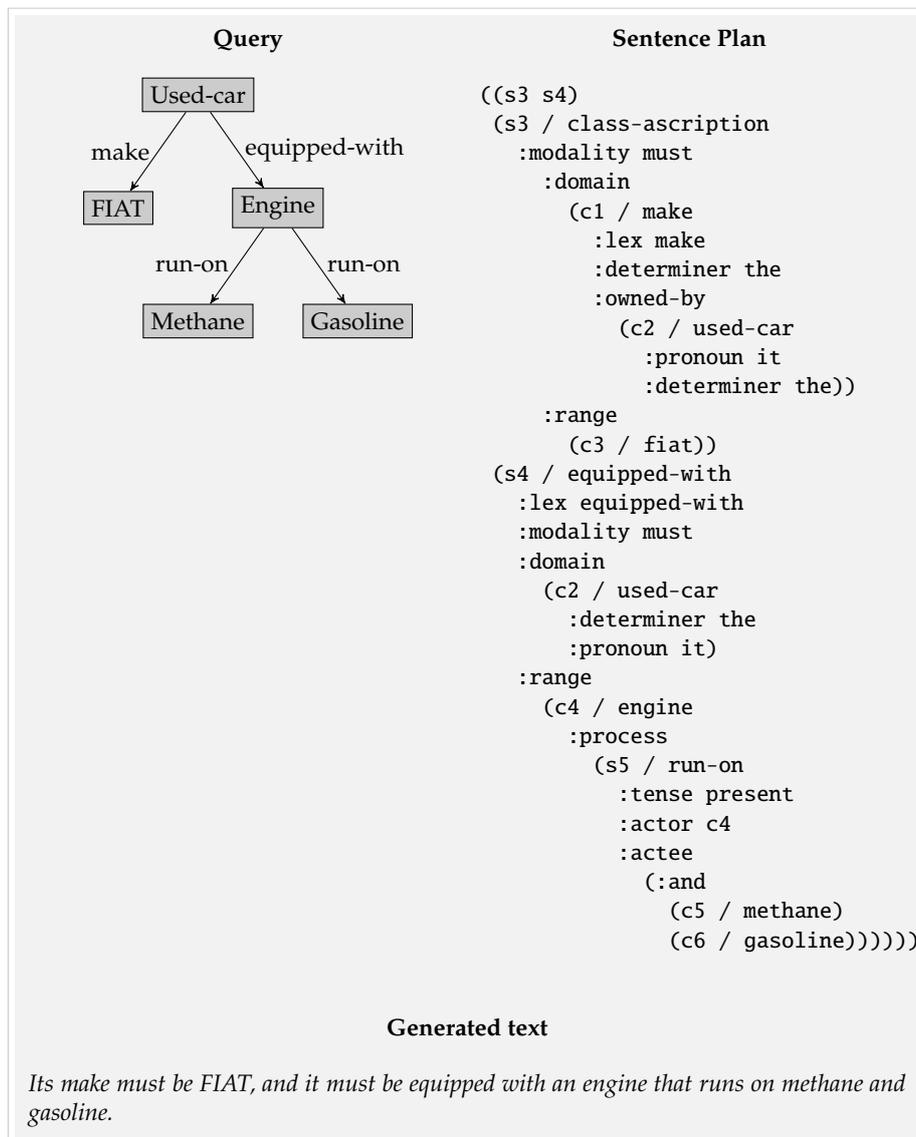
These three concepts are represented in the sentence plan by three variables  $c1$ ,  $c2$ , and  $c3$ ;  $s1$  instead, is the variable representing the *relational process* we use in order to verbalize our input query in a descriptive way. The process is a *class-ascription*, one of the process types defined in the *Merged Upper Model* (see Section 6.3), a general task- and domain-independent linguistically-motivated ontology used for mediating between domain knowledge and the linguistic realizer. A *class-ascription* process must have at least two participants, which are called *:domain* and *:range*: The domain is the first concept (*Used-car*), and the range is the conjunction of all other concepts (in this case *Off-roader* and *Non-smoker-car*).  $\kappa\text{PML}$  generates the class ascription as a copula that relates domain and range as subsets i.e. the used car we are looking for is contained in the intersection of the sets of all off-roaders and non-smoker cars. We also added the *:modality* property to the class ascription process, in order to emphasize that this is a query expressing user requirements.

**Example 2** Here the query is composed by three concepts (*Used-car*, *Air-conditioning*, and *Central-locking*) and two instances of the same role (*equipped-with*).



The derived sentence plan contains a process named `equipped-with` which is subsumed by the more general Upper-Model (UM) concept called `generalized-possession`. The participants are the used car as `:domain` and both air-conditioning and central-locking as `:range`. We decided to pronominalize the subject of the sentence.

**Example 3** We show here a more complex query, containing five concepts and three roles.



The sentence plan is made up of two main coordinate clauses, s3 and s4, which are associated to two processes: a class-ascription and equipped-with (seen in the previous example). The latter contains a further process (run-on) that gives additional information about the engine's fuel (methane and gasoline): This sub-process is realized as a relative clause (*... that runs on methane and gasoline*). Since run-on is subsumed by the UM-process *dispositive-material-action*, the participants of this process have to be named :actor and :actee. We also want to remark the use of a possessive (*its*) and a definite pronoun (*it*) referring to used-car, along with the relative pronoun (*that*) referring to the engine, automatically generated by the realizer as subject of the sub-process.

We proceed now formally by describing how to map each one of the templates we listed in Table 3 into a corresponding text plan. Indexes of concepts in that table are equal for concepts on the same level, while here indexes are

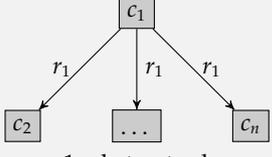
numbered differently for different concepts. For each template, we show the corresponding generic query, its linearization and the generated sentence plan.

We draw the attention of the reader to the fact that all concepts included in the following sentence plans may seem confusing because of a name duplication. When we write  $(c_i / c_i)$ , the first  $c_i$  is a variable, the second one is the name of the concept assigned to the variable, also called *type* as we will see in Section 6.4.

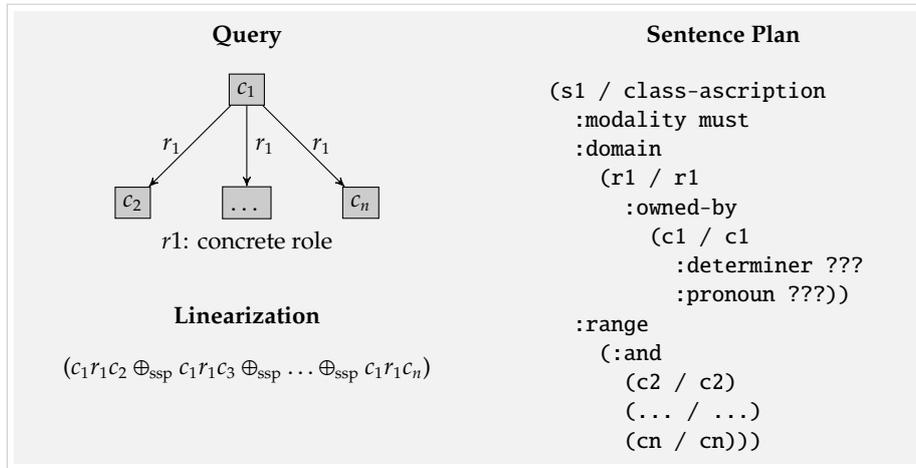
**Template 1.**  $r_{is}$  is the “compatibility relation” we introduced on page 33 and that is mapped to a class-ascription in the sentence plan..

Query	Sentence Plan
$c_1 \sqcap c_2 \sqcap \dots \sqcap c_n$	(s1 / class-ascription :modality must :domain (c1 / c1 :determiner ??? :pronoun ???) :range (:and (c2 / c2) (... / ...) (cn / cn)))
<b>Linearization</b> $(c_1 r_{is} c_2 \oplus_{ssp} \dots \oplus_{ssp} c_1 r_{is} c_n)$	

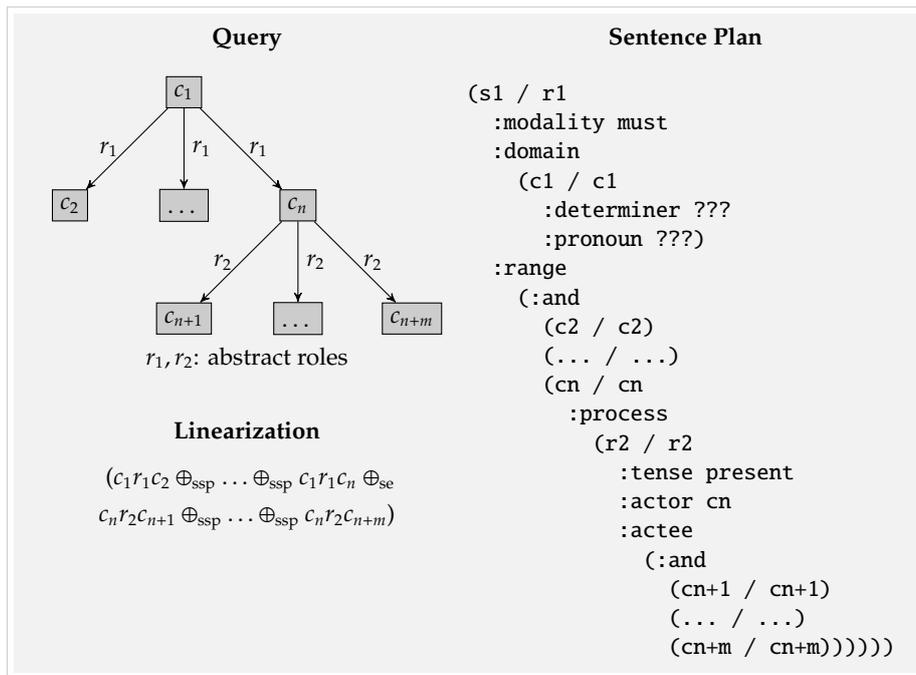
**Template 2.**  $r_1$  is an *abstract role* in the domain ontology:

Query	Sentence Plan
 <p><math>r_1</math>: abstract role</p>	(s1 / r1 :modality must :domain (c1 / c1 :determiner ??? :pronoun ???) :range (:and (c2 / c2) (... / ...) (cn / cn)))
<b>Linearization</b> $(c_1 r_1 c_2 \oplus_{ssp} c_1 r_1 c_3 \oplus_{ssp} \dots \oplus_{ssp} c_1 r_1 c_n)$	

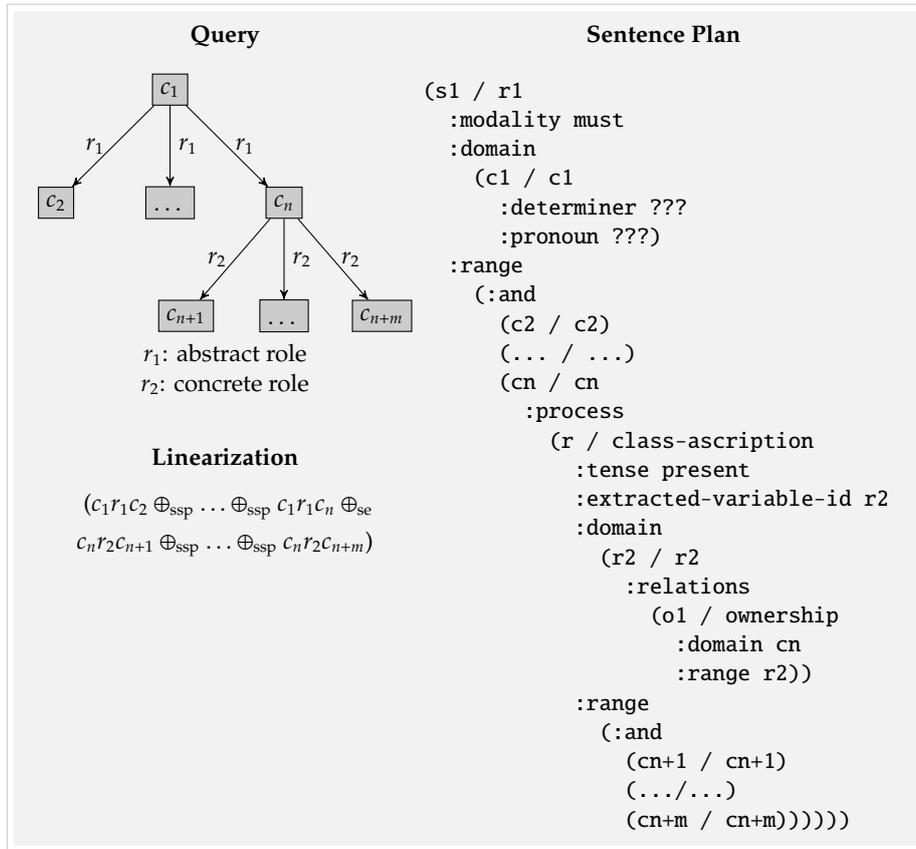
or, if  $r_1$  is a *concrete role* in the domain ontology:



**Template 3.** If  $r_1$  and  $r_2$  are both abstract roles:

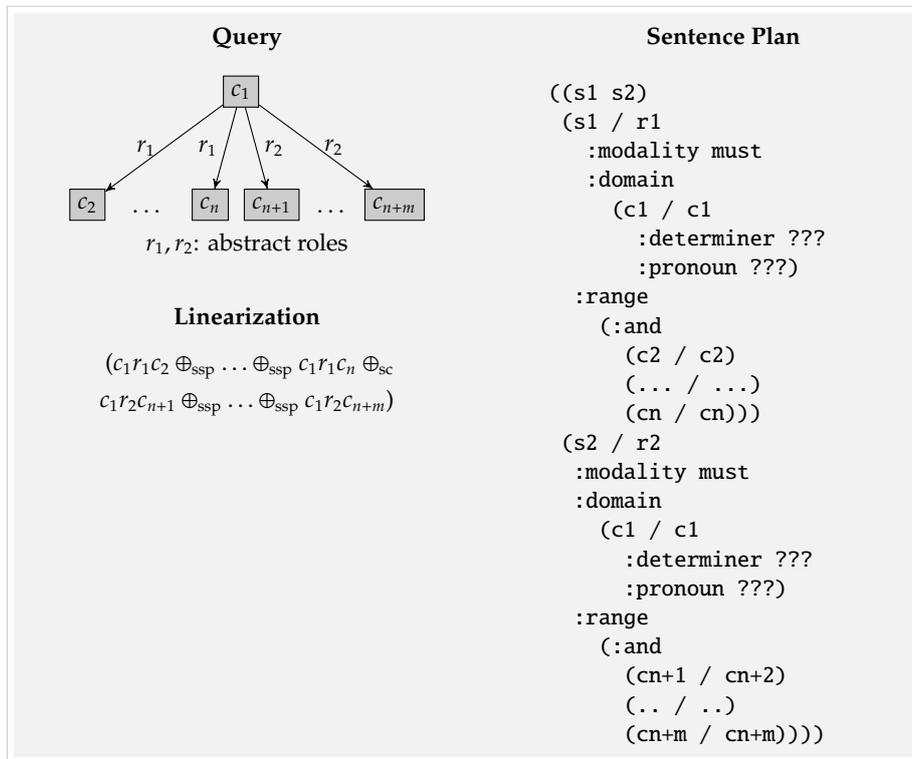


otherwise, if  $r_2$  is a concrete role we have:



The attribute `:process` in the SPL code above and in the sentence plan of Example 3 is used to produce a complex modification of the preceding concept ( $c_n$ ), modification that will be rendered as a relative clause.

**Template 4.** Here we have four possibilities, depending on the fact that  $r_1$  and  $r_2$  can either be abstract or concrete roles. We start with the case that  $r_1$  and  $r_2$  are abstract roles.



If either one of  $r_1$  or  $r_2$  (or both) is a concrete role (say  $r_y$ ), the previous sentence plan is no longer valid. We need to replace the sentence plan chunk containing  $r_y$  (left box below) with the one on the right.

```

(sx / ry
 :modality must
 :domain
 (c1 / c1
  :determiner ???
  :pronoun ???))

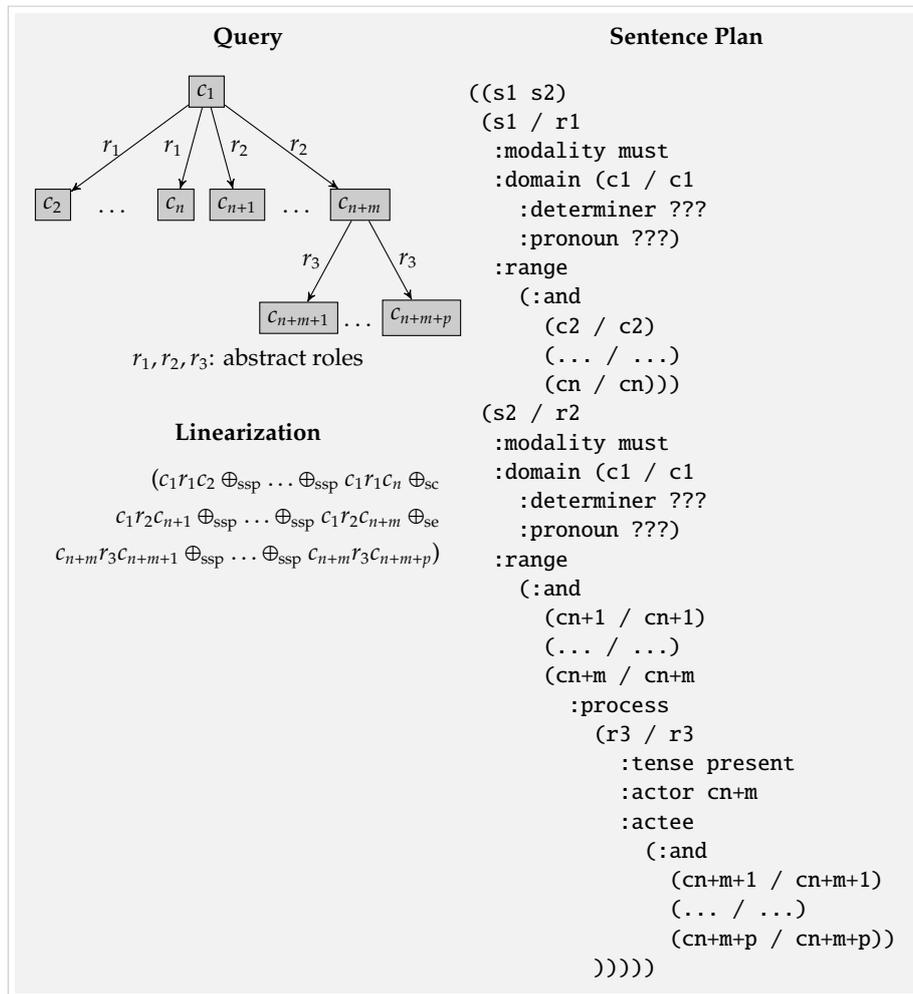
```

```

(sx / class-ascription
 :modality must
 :domain
 (ry / ry
  :owned-by
  (c1 / c1
   :determiner ???
   :pronoun ???))

```

**Template 5.** If  $r_1, r_2,$  and  $r_3$  are abstract roles:



If  $r_1$  is a concrete role we should substitute in the previous plan the chunk reported in the left box below with the one on the right box.

```

(s1 / r1
 :modality must
 :domain (c1 / c1
  :determiner ???
  :pronoun ???)

```

```

(s1 / class-ascription
 :modality must
 :domain
 (r1 / r1
  :owned-by
  (c1 / c1
   :determiner ???
   :pronoun ???))

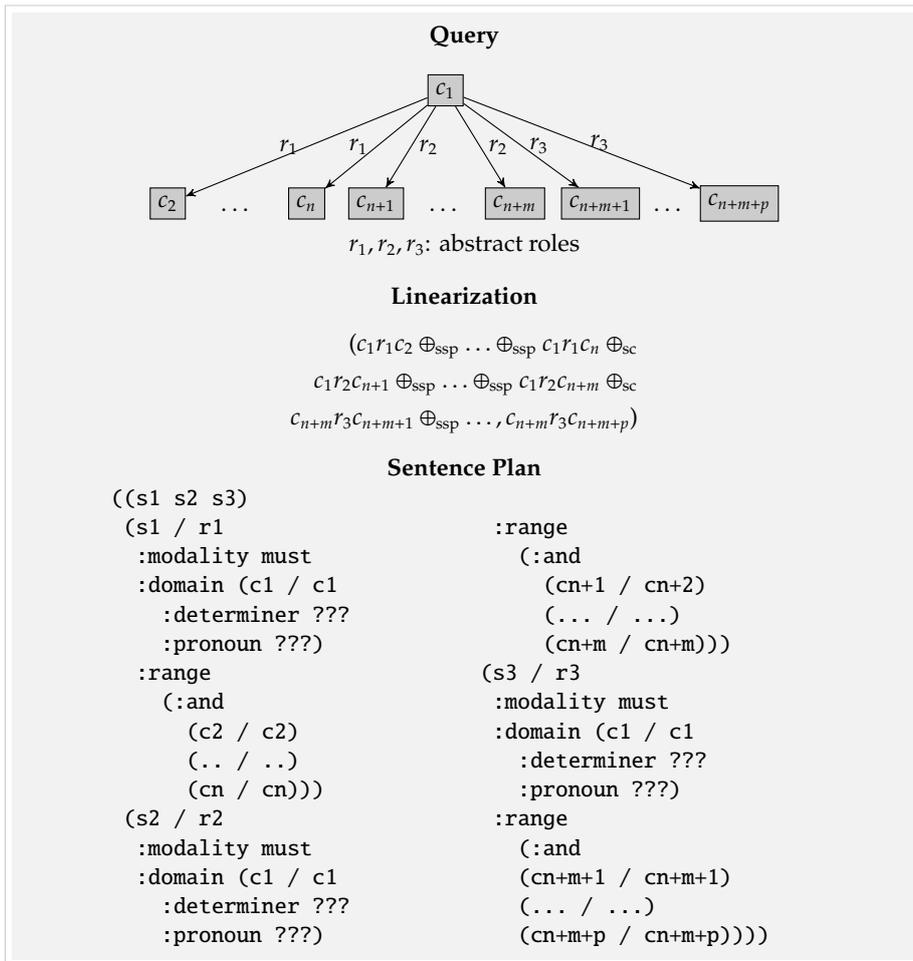
```

Since  $r_2$  cannot be a concrete role (otherwise  $c_{n+m}$  would be a concrete data type, which is not possible because it should be a leaf), the last variant to this sentence plan is that  $r_3$  is a concrete role. The substitution we perform in this case is:

```
(r3 / r3
:tense present
:actor cn
:actee
(:and
(cn+m+1 / cn+m+1)
(... / ...)
(cn+m+p / cn+m+p)))
```

```
(s3 / class-ascription
:tense present
:extracted-variable-id r3
:domain
(r3 / r3
:relations
(o1 / ownership
:domain cn
:range r3))
:range
(:and
(cn+m+1 / cn+m+1)
(... / ...)
(cn+m+p / cn+m+p)))
```

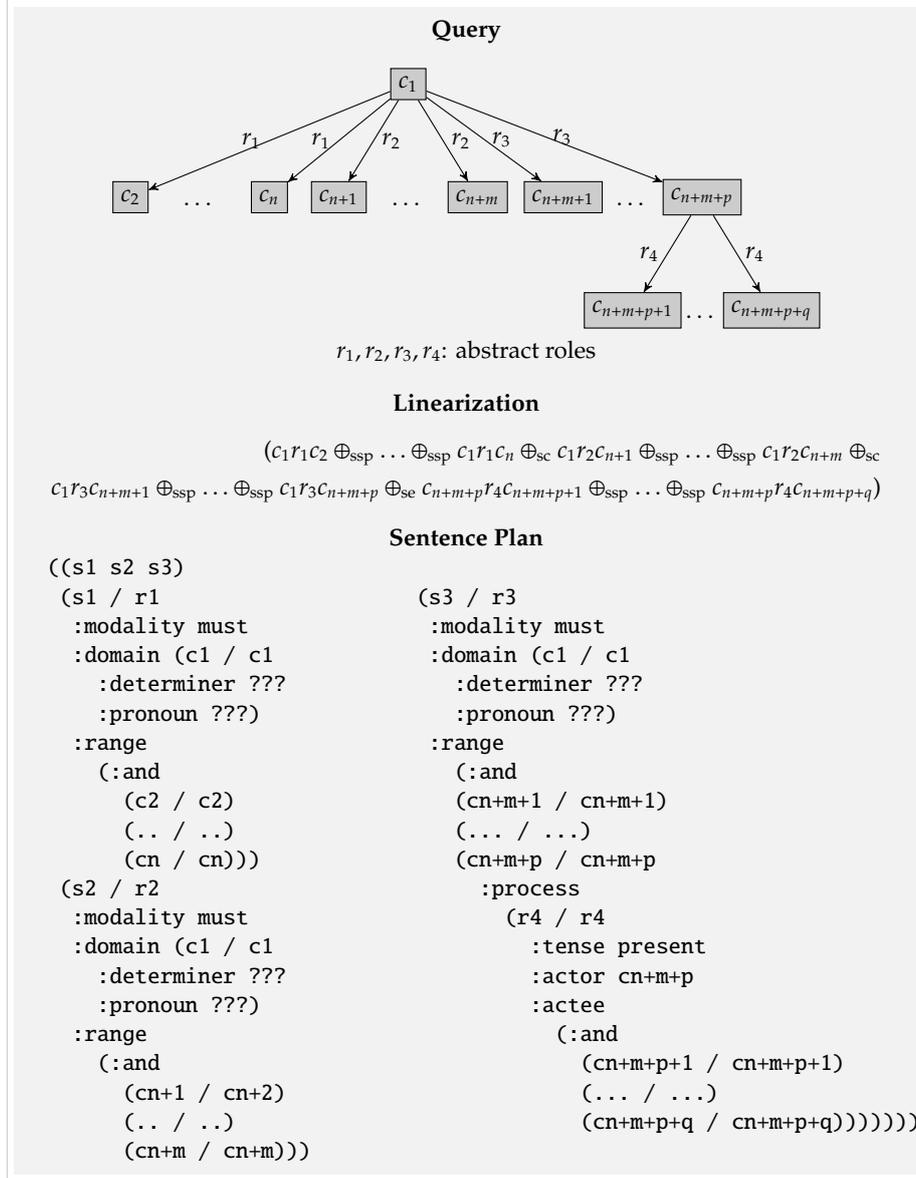
**Template 6.** If  $r_1, r_2,$  and  $r_3$  are abstract roles:



Otherwise, if any of  $r_1, r_2,$  or  $r_3$  is a concrete role, we perform a substitution

as we explained for Template 4.

**Template 7.** If  $r_1, r_2, r_3,$  and  $r_4$  are abstract roles:



If any of  $r_1, r_2,$  or  $r_4$  is a concrete role, we perform a substitution as explained for Template 4.

**Fine tuning of the sentence plans** Given the text plan augmented with aggregation information (i.e. the matching templates from Table 3) and the referring expression types (Table 5) associated to each concept, obtaining the sentence plan(s) is just a matter of mapping the list of aggregation patterns into the corresponding SPL chunks as specified above.

It was not mentioned above that the sentence plan(s) must be preceded by an introductory one, which declares what the user is looking for (root concept, say  $c_1$ ), at least in our setting where the complex concept description represents a user query.

introductory  
sentence plan

One possible introductory sentence could be “I’m looking for  $c_1$ ”, where  $c_1$  is the root concept of the query, and the corresponding sentence plan would be the following:

```
(s0 / look-for
  :actor speaker
  :actee (c1 / c1)
  :tense present-continuous
)
```

The type `look-for` is subsumed by the UM-process `dispositive-material-action`, the participants of this process are `:actor` and `:actee`; the actor in this case is of type `speaker` and will be realized as a first person singular (I), and the actee will be the  $c_1$ . The verb describing the process `look-for` will be rendered as a present continuous.

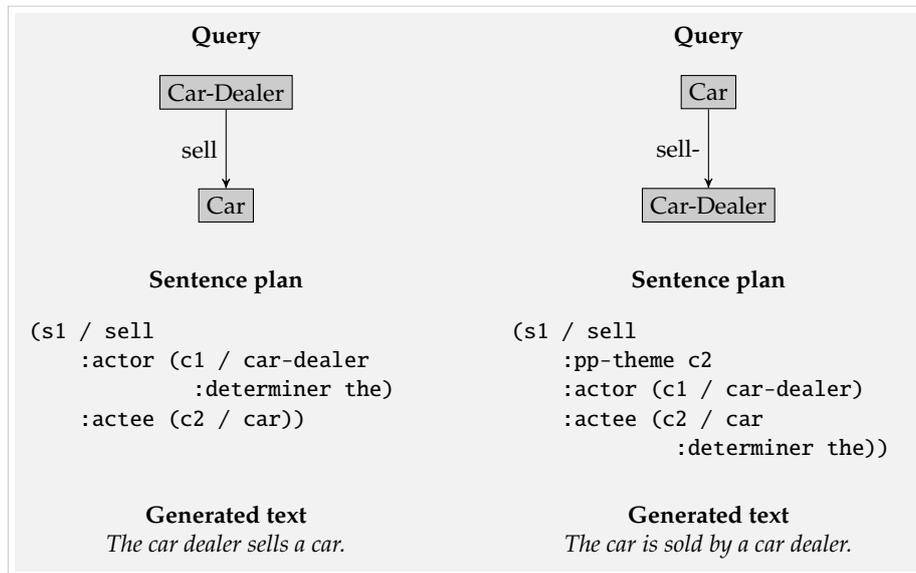
With this initial SPL chunk, we have all needed sentence plans to be passed to the linguistic realizer. They need to be finalized though.

First of all we need to check if any relation (abstract role) has to be realized as a passive voice verb. In order to test this, we need a way to directly recognize from the role name if the domain concept (subject) will act as source (actor) or recipient (actee) of the action represented by the role. We will call the former kind of role *active role* and the latter as *passive role*. We can solve this in two ways. The first solution is adopting a simple naming convention that helps us recognize passive roles, which is a *minus sign* (–) put as suffix of the corresponding active role. E.g. if the active role is `sell`, the passive role will be `sell-`. Therefore, whenever a query contains a passive role, its minus sign is deleted, and the active role is used in the sentence plan with an additional SPL line (`:pp-theme`) specifying that the grammatical subject is the recipient of the action denoted by the verb. The second solution consists in avoiding naming conventions and building instead a correspondence map between passive roles and the respective active roles, as for example:

passive voice

Passive role	Active role
<code>soldBy</code>	<code>sell</code>
<code>madeBy</code>	<code>make</code>
...	...

We show below two sentence plans involving an active role and the corresponding passive role.



A further adjustment of the generated sentence plans is the correct assignment of values to the parameters `:determiner` and `:pronoun` according to the output of Algorithm 9 together with Table 5. For each concept, the output of the algorithm can be one of R-INP, R-DNP, R-DP, R-RPS, R-PP, R-RPP. For each one we detail now what the effects on the sentence plan (hereinafter SP) are.

referring  
expressions

**R-INP** in this case nothing needs to be done because if no `:determiner` attribute is specified in the SP, the linguistic realizer automatically assigns an indefinite article (*a, an*) if the entity is a *common countable noun*. If the entity is a *proper noun* or a *mass noun*, no indefinite article will be used; the same happens for adjectives and instances of any concrete datatype (a generic string, a number, a range of numbers, a date etc.).

**R-DNP** this triggers the assignment of the article *the* to the `:determiner` attribute in the SP (only for common nouns);

**R-DP** here we must consider first if the concept in question is *human* or *non-human*. This information must be available in the ontology, e.g. under the form of concepts as `Non-Human-Entity` and `Human-Entity` directly or indirectly subsuming the given concept. If human, we need to check if it's either `Male` or `Female` using the pronoun *he* or *she* respectively, or *they* if undefined. For non-human entities we use *it*. In the SP we assign one of `{it, he, she, they}` to the attribute `:pronoun`;

**R-RPS** nothing needs to be done here, because the SP structure will already lead the linguistic realizer to render the concept as a relative pronoun (see e.g. the process that involves actor  $c_n$  in Template 3;

**R-PP** the same considerations made for **R-DP** are valid here; the pronoun values of the set `{it, he, she, they}` are used and associated via the `:pronoun` attribute to the owner of a given concrete role; the linguistic realizer translates then the definite pronoun into the right possessive pronoun;

**R-RPP** for the generation of relative pronouns as possessives, nothing needs to be added to the SP.

The following last consideration shortly discusses the inclusion of a modal verb in the clauses, needed to emphasize the requirements of the user in terms of relations among concepts and their attributes in a query. The modal we already introduced before is *must* as e.g. in the sentence *The car must be an off-roader and a non-smoker car*. Naively putting it in each sentence would be quite annoying for the same user rereading the query, but on the other side this would be the right way of referring to an object we are looking for and precisely describing. modal verb

An option would be the one of rendering the query without modal auxiliaries (in our case only *must*), and the user describes the object how it *is* instead of how it *must be*. The introductory sentence could be extended e.g. as: *I'm looking for a THING that is described as follows*, where THING stands for the starting concept chosen by the user. The SP generating this sentence is:

```
(s0 / look-for
  :tense present-continuous
  :actor speaker
  :actee
    (c1 / THING
      :process
        (s1 / be-described-as
          :tense present
          :actor c1
          :actee
            (c2 / template
              :pattern "follows")))))
```

**A Complete SPL Example** Finally we present the sentence plans we obtain from the query of Figure 9 with the outputs from the discourse planning, sentence aggregation, and referring expression generation phases.

```
(s0 / look-for
  :actor speaker
  :actee (c1 / car)
  :tense present-continuous
)
```

$S_0$ : *I'm looking for a car.*

```

(s1 / class-ascription
 :domain
  (c1 / car
   :determiner the
   :pronoun it)
 :range
  (:and
   (c1.1 / off-roader)
   (c1.2 / non-smoker-car))
 )

```

*S<sub>1</sub>: It is an off-roader and a non-smoker car.*

```

((s2.1 s2.2)
 (s2.1 / make
  :pp-theme c2
  :actor (c4 / land-rover)
  :actee (c1 / car
   :determiner the
   :pronoun it))
 (s2.2 / be-equipped-with
  :domain
  (c1 / car
   :determiner the
   :pronoun it)
  :range (c2 / engine
   :process
   (s19 / run-on
    :tense present
    :actor c2
    :actee (:and
     (c3 / diesel)
     (c6 / electrical-power))))))

```

*S<sub>2</sub>: It is made by Land Rover, and it is equipped with an engine that runs on diesel and electrical power.*

```

(s3 / sell
 :pp-theme (c1 / car)
 :actor (c5 / car-dealer
 :process
 (r / class-ascription
 :tense present
 :extracted-variable-id r7
 :domain (r7 / name
 :relations
 (o1 / ownership
 :domain c5
 :range r7 ))
 :range (c8 / template
 :pattern "[...]"))))
 :actee (c1 / car
 :determiner the
 :pronoun it))

```

*S<sub>3</sub>: The car is sold by a car dealer whose name is [...].*

```

(s4 / situated-in-city
 :domain (c5 / car-dealer
 :determiner the
 :pronoun they
 :range (c7 / city
 :process
 (r8 / located-in-country
 :tense present
 :domain c7
 :range (c9 / italy))))))

```

*S<sub>4</sub>: They are situated in a city that is located in Italy.*

```

(s5 / located-in-province-of
 :domain (c7 / city
 :determiner the
 :pronoun it)
 :range (c10 / trento))

```

*S<sub>5</sub>: It is located in Trento.*

```
(s6 / class-ascription
  :domain
  (r10 / phone-number
    :owned-by (c5 / car-dealer
               :determiner the))
  :range
  (c11 / template
    :pattern "[...]"))
```

*S<sub>6</sub>: The car dealer's phone number is [...].*

Feeding the six sentence plans presented above into a linguistic realizer supporting SPL (as e.g. the KPML system presented in the coming section), the complete text output would read as follows:

*I'm looking for a car. It is an off-roader and a non-smoker car. It is made by Land Rover, and it is equipped with an engine that runs on diesel and electrical power. The car is sold by a car dealer whose name is [...]. They are situated in a city that is located in Italy. It is located in the province of Trento. The car dealer's phone number is [...].*

## 5 Linguistic Realization

Linguistic realization is the last operation of the NLG-pipeline we have described so far. The task of a linguistic realizer is to convert sentence-sized chunks of a suitable input representation (sentence plan) into grammatically correct sentences.

### 5.1 Approaches to LR

Four main approaches to text realization are available [Hovy, 1997], differing in terms of sophistication/expressive power and flexibility. They are listed below, ordered from the simplest (and less flexible) to the most sophisticated ones:

- canned text systems,
- template systems,
- phrase-based systems,
- feature-based systems.

We will describe in turn each one of these categories.

#### 5.1.1 Canned Text Systems

Whenever we want to generate a piece of text for a very specific purpose, without the need of modifying it according to some parameters, canned text is the easiest solution. It has been used by almost every application to convey a message (warning, error, help, etc.) to the user, a message which is simply associated with a given code produced by an application event. No syntactic or morphological process is involved, except, in some cases, capitalizing the first word in the sentence, and putting a full stop at the end.

### 5.1.2 Template-based Systems

Slightly more sophisticated, these systems provide template texts containing a certain number of placeholders, which at runtime will be substituted with strings depending on the context (a title, a name, an address, some numbers, etc.). One typical example is represented by mail-merge applications, where the same letter with a few variations (receiver, salutation, closing, ...) needs to be created in multiple copies for different receivers.

### 5.1.3 Phrase-based Systems

In these systems, templates are more general, and resemble phrase structure grammar rules. They represent the various typologies of phrases we have in natural language (noun phrases, verb phrases, etc.) along with a set of rules specifying how phrases can be combined together to form grammatical sentences. E.g. we could have a pattern like [subject verb object] where each one of its components can be further decomposed into one of other possible phrasal patterns as [subject] → [determiner adjectives head-noun modifiers]. The generation process starts with a top-level sentence pattern matching the sentence plan, and stops when all pattern constituents have been replaced by one or more words.

The phrase-based approach is quite flexible in comparison to the ones seen before, and it is rather simple to implement such a system for a grammar of limited size; beyond a certain limit though, it is hard to keep track of all phrasal interrelationships in order to avoid wrong phrase expansions.

### 5.1.4 Feature-based Systems

These represent the highest level of sophistication and flexibility available for the generation of sentences. Here every possible alternative for expressing a sentence or part of it can be chosen by means of features: we can say if a sentence is *positive* or *negative*, if it is *declarative*, *imperative*, or a *question*, which are the tenses used in its clauses, etc. Generation in this case is accomplished by incrementally collecting features for each part of the input sentence plan until the sentence is complete: this can be done either by traversing a *feature selection network* (see Sec. 6.1) or via *unification* ([Kay, 1979]).

The strength of a feature-based approach is that any distinction in language can be encoded as a feature in the system. On the other side, cons of this approach are that also in these systems—as in the previous ones—maintenance of feature interrelationships tends to be quite hard; moreover some authors [McRoy *et al.*, 2001] report that since quite often the entire grammar needs to be traversed, such systems had previously tended to be too slow for real-time applications, but this is no longer an issue nowadays.

## 5.2 The Linguistic Realizer of Choice

Among the available feature-based linguistic realizers, KPML is the system we employed and we just provide a short description of it in this section, leaving all details for Section 6.

KPML (KOMET-Pennman-multilingual) is a grammar development environment from the University of Bremen [Bateman, 1997a]. KPML is a complex application, well known for extensive multilingual systemic-functional grammar (SFG) development and maintenance as well as for NL generation. For the sake of preciseness, as described in [Bateman, 1997b], the intended purposes of KPML are:

- to offer generation projects large-scale, general linguistic resources (at the time of writing available resources include English, Chinese, Czech, Greek, Japanese, Russian, German, and Spanish in varying stages of development);
- to offer generation projects an engine for using such resources for generation;
- to encourage the development of similarly structured resources for languages where they do not already exist;
- to provide optimal user-support for undertaking such development and refining general resources to specific needs;
- to minimize the overhead of providing text in multiple languages;
- to encourage contrastive functional linguistic work;
- to raise awareness and acceptance of text generation as a useful endeavor.

KPML can be used as fully featured grammar development environment, but it is also available as a simple blackbox linguistic realizer. The environment offered by the system takes over and extends the functionality of its predecessor, the Penman text generation system [Mann, 1983a; 1983b] outperforming it in terms of ease of use, development support, and multilingual design.

The input required by KPML is an annotated semantic specification (sentence plan) expressed using the *Sentence Plan Language* (SPL). Our sentence planner, as shown above, adopts this language which will be formally described in Sec. 6.4.

## 6 Linguistic Realization with Systemic Functional Grammar

In the previous section we introduced several approaches to linguistic realization directing our attention towards one in particular: Feature-based realization. In this section we refine our choice a little further describing linguistic realization done by means of a famous and fascinating linguistic theory, Systemic Functional Linguistics (SFL), which was leveraged for the purpose of natural language generation giving rise to what is called *computational SFL*. We present Systemic Functional Grammar (SFG) and a computational implementation called the Nigel systemic grammar of English. We will see what the *Upper Model*, a linguistic ontology is used for, and what is the input specification to the chosen realizer we chose (KPML). Finally we will see how KPML really works and how we are using it for our purpose.

## 6.1 Systemic Functional Grammar

### 6.1.1 History

Systemic Functional Grammar (SFG) is a grammar model and major influential linguistic theory developed by Michael Alexander Kirkwood (M. A. K.) Halliday. It grew out of the work of John Rupert Firth, a British linguist who influenced a whole generation of linguists for more than twenty years at the University of London. Firth was an important figure in the foundation of linguistics as an autonomous discipline in Britain, and the popularity of his ideas among contemporaries gave rise to what was known as the 'London School' of linguistics. Among Firth's students, the so-called *neo-Firthians* were exemplified by Michael Halliday, Professor of General Linguistics in the University of London from 1965 until 1970 when he moved to Australia, establishing the department of linguistics at the University of Sydney. Through his teaching there, SFL has spread to a number of institutions throughout Australia, and around the world.

### 6.1.2 Theory

Systemic-functional grammar is concerned primarily with the *choices* that are made available to speakers of a language by their grammatical systems. These choices are assumed to be meaningful and relate speakers' intentions to the concrete forms of a language.

Language is considered a social resource by means of which speakers and hearers act meaningfully. Meanings in systemic functional grammar are divided into three broad areas, called *metafunctions*: the *ideational*, the *interpersonal* and the *textual*, as extensively described in the SFL literature, in particular [Halliday and Matthiessen, 2004].

- The *ideational* is grammar for representing the world. That is, the propositional content, which is concerned with ideation providing the speaker with the resources for interpreting and representing reality. It is divided into two subtypes, the *experiential* and the *logical* metafunctions: The former is reflected in terms of configurations of processes and participants. We could name e.g. the TRANSITIVITY structure of the clause, which describes what in other theories are known as semantic relations. The experiential part of the ideational metafunction also includes systems for circumstantials, types of prepositional phrase, tense, noun-types, etc. The logical part, instead, is the mode for creating various kinds of complexes that are hypotactically or paratactically related.
- The *interpersonal* is grammar for enacting social relationships such as asking, requests, asserting control, or ordering. Thus the interpersonal metafunction is very much about interaction between human beings, society and culture.
- Finally, the *textual* is grammar for binding linguistic elements together into broader texts (via pronominalizations, grammatical topicalization, thematization, expressing the newsworthiness of information, etc.), or more simply, the rhetorical structure of a text. What is a subordinate

clause? What is an independent clause? These are the kinds of questions that deal with the textual element of meaning.

Systemic functional grammar deals with all of these areas of meaning equally and within the grammatical system itself.

From a higher level perspective, as clearly explained in [Teich, 1999, pages 8–9], sFL view of language rests upon four main considerations:

- language is behaviour potential, realized by systems that support the theoretical notion of choice;
- language construes meaning which is realized by stratification (phonology, grammar, semantics), represented in sFL as paradigmatically organized resources;
- language is multifunctional, where functional diversification is represented by the metafunctions described above;
- using language is choice in the potential and ultimately actualization of the potential, by means of realization statements.

In the following subsections we will see in more detail what a system and a system network are, along with an explanation on how to specify linguistic structure by means of realization.

**The System Network** A *system network* is a directed graph whose nodes are choice points called *systems*. Each system consists of *entry conditions* and *output features*. A small section of systemic network for the English grammar is shown in Figure 12, where system names are capitalized.

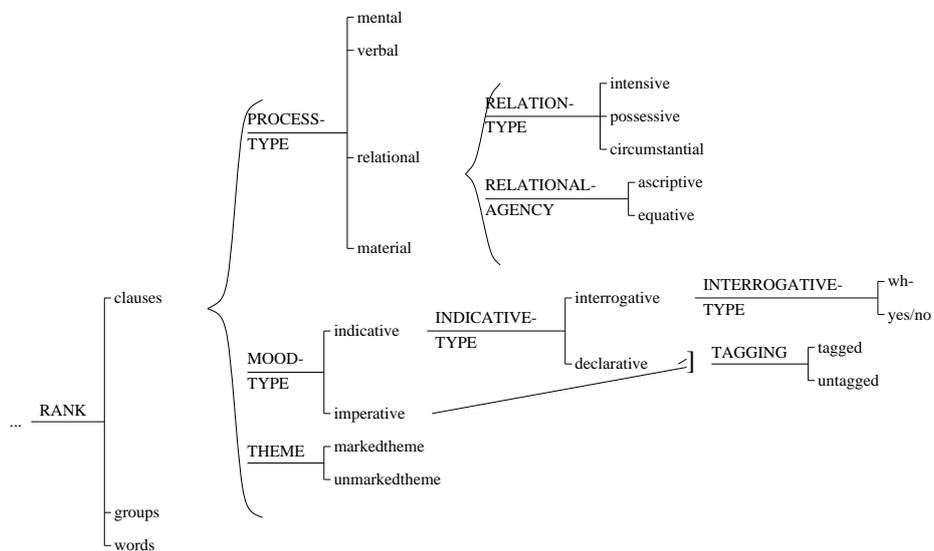


Figure 12: Example of system network fragment

The “MOOD TYPE” system e.g. has an entry condition which is “clause”, and two alternative output features which are “indicative” and “imperative”.

Entry conditions can be conjunctions or disjunctions of output features of other systems. The “TAGGING” system for example has a disjunctive entry condition, which can be either the “declarative” or “imperative” feature. There can be *simultaneous systems* that share entry conditions, such as “PROCESS TYPE” and “MOOD TYPE”; this means that both are relevant in the paradigmatic context described by the entry condition “clause” and both must be entered as soon as the system “RANK” outputs “clause”.

Connections among systems define a partial ordering that spans, if we consider the graphical representation, from least *delicate* (most general) systems on the left to most delicate (most specific) systems on the right. We have an *incremental description refinement* as discussed in [Mellish, 1988], a scale of delicacy representing a left-to-right dimension. An interesting aspect is that paradigmatic choices in systems take place not only between grammatical alternatives, but also between lexical alternatives. In fact Halliday introduces the term *lexico-grammar* to include both of them, meaning that there is no clear division between grammar and lexicon, and if on the left part of the network we have grammatical choices, towards the right side of it lexical choices take place. This is summarized in Halliday’s expression of *lexis as the most delicate grammar*.

We still need to see how systems are related to the functional side of SFG, in particular with metafunctions. The relation is that each system pertains to one and only one metafunction. Moreover, systems of the same metafunction are strictly connected, in a measure that they are largely independent from systems of other metafunctions. If we refer to Figure 12, “PROCESS-TYPE” and the systems depending on it are in the “TRANSITIVITY” region of the ideational metafunction; “MOOD TYPE” and its successors are in the “MOOD” region<sup>3</sup> of the interpersonal metafunction, while “THEME” and other systems connected to it are in the “MOOD” region of the textual metafunction. Table 6 shows the main systems in SFG according to some high-level entry conditions and the three metafunction.

	<b>ideational</b>	<b>interpersonal</b>	<b>textual</b>
<b>clause</b>	TRANSITIVITY	MOOD	THEME
<b>verbal group</b>	TENSE	MODALITY	VOICE
<b>nominal group</b>	MODIFICATION	PERSON	DETERMINATION

Table 6: Main systems in SFG

**Specifications of linguistic structure** The way syntactic structure is created in SFG is by means of *realization statements* which are associated with the output features of systems, and show how the paradigmatic choices in the systems are expressed as syntagmatic chains in the language structures. In Figure 13 we show a system network fragment augmented with realization statements.

<sup>3</sup>*Regions* are groups of systems within the same metafunction, possessing strong intra-region dependencies, and weak inter-region dependencies, creating a modularity that is beneficial for grammar design, maintenance and development.

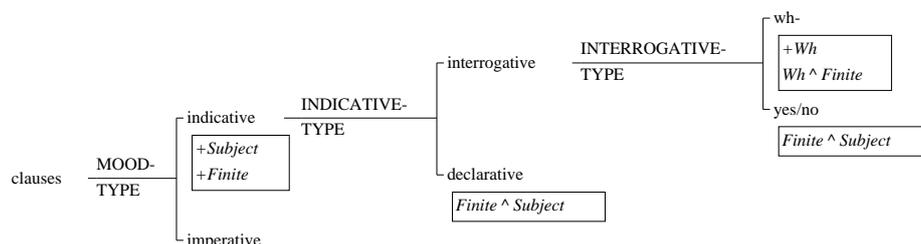


Figure 13: Example of system network fragment with realization statements

The “indicative” feature e.g. embeds two realization statements, “+Subject” and “+Finite” which are *insertion* realizations. The “yes/no” feature instead has just one, “Subject  $\wedge$  Finite”, which is an *ordering* realization. Table 7 summarizes the realization statements of sfg.

Name	Notation and example	Description
<b>insert</b>	+Subject	this statement requires the presence of this function as constituent
<b>order</b>	Subject $\wedge$ Finite	this requires that the two functions must be ordered one after the other
<b>conflate</b>	Subject / Agent	requires that the two functions are realized by the same element of structure
<b>expand</b>	Mood (Finite)	the first function is expanded to have the one in brackets as constituent
<b>preselect</b>	Subject : singular	this constrains the realization of the function to display the given feature

Table 7: Realization statements used in sfg

We terminate with a simplified example (Figure 14) of the kind of information that is specified in an sfg syntagmatic unit.

<i>In your car</i>	<i>Paul</i>	<i>you</i>	<i>will</i>	<i>find</i>	<i>a navigation system</i>	
Theme		Rheme				<b>textual</b>
		Mood		Residue		<b>interpersonal</b>
	Vocative	Subject	Finite			
Locative		Actor	Process	Complement		<b>ideational</b>

Figure 14: Simplified example of metafunctional layering

## 6.2 The Nigel systemic grammar of English

Nigel represents the biggest freely available computational systemic functional grammar for English available to-date. Nigel has been under development

since the early 1980s [Matthiessen, 1981; 1983; Mann and Matthiessen, 1983], when it was used within the Penman project for English generation. It was mainly developed by Christian Matthiessen on the foundation of work by Michael Halliday. Since then many people have contributed to various parts of its coverage.

The latest version consists of around 765 systems<sup>4</sup>, where the first one to be entered is the “RANK” system reported below, whose output features are the items of the *rank scale*<sup>5</sup>: clauses, group-phrases, words or morphemes.

rank scale

```
(SYSTEM
  :NAME RANK
  :INPUTS START
  :OUTPUTS
    ((0.2 CLAUSES)
     (0.2 GROUPS-PHRASES)
     (0.2 WORDS
      (INSERT STEM)
      (PRESELECT STEM MORPHEMES))
     (0.2 MORPHEMES
      (INSERT HEAD)))
  :CHOOSER RANK-CHOOSER
  :REGION RANKING
  :METAFUNCTION LOGICAL
)
```

At word and morpheme level, the Nigel grammar does not provide a unified lexicogrammar of SFG as in the theory; lexis and morphology are treated apart in an external lexicon. At the clause level the grammar can generate clause complexes of two clauses in paratactic or hypotactic relation. In order to generate, the system network is traversed starting from the “RANK” system; the rule is that whenever an output feature is chosen, the next step is to collect all systems having the same entry conditions as the preceding output feature, and to enter each one of them on turn in random order, provided all other entry conditions are satisfied. Every time an output feature is chosen, the realization statements attached to it are immediately executed, except the ordering realizations which are collected and executed later.

The choice among output features is done by means of choosers and inquiries, an explicit formalization developed by William C. Mann under the name of *inquiry semantics* or *chooser/inquiry interface* [Mann, 1983a]. Each system with more than one output feature specifies a *chooser*, a small “choice expert” that knows how to make appropriate choices among the grammatical features available. This is done by traversing a decision tree from the root to one of the leaf nodes which represents the chosen feature. *Inquiries* are oracles which can be relied on to motivate grammatical alternations for the current communicative goals being pursued. Figure 15 shows the same network of

<sup>4</sup>The total count of systems includes 324 *gates* which are symplified system having only one output feature.

<sup>5</sup>The rank scale defines the typologies of linguistic units used in the grammar. Units are hierarchically ordered by rank according to their constituency relation: higher-ranking units are built with units of the next lower rank. Units categorized under the lowest rank cannot be decomposed further. The English grammatical rank scale consists of *clause*, *group/phrase*, *word*, and *morpheme*.

figure 13 augmented with choosers.

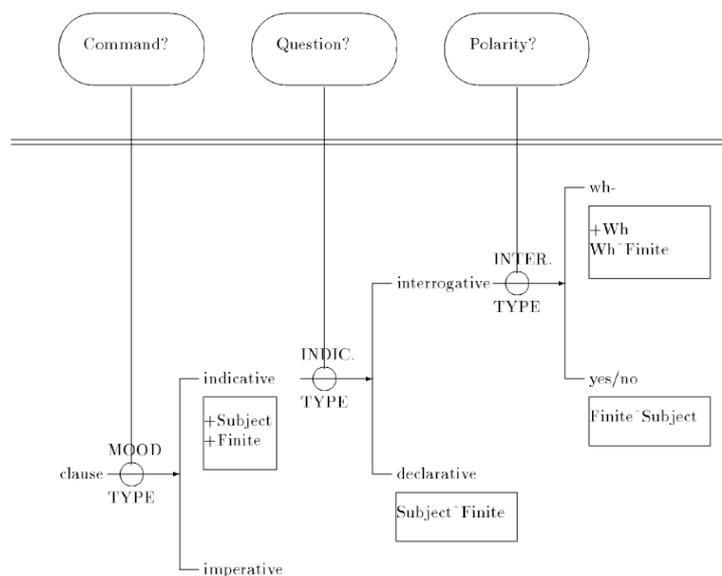


Figure 15: Example of system network with choosers

### 6.3 The Upper Model

The Upper Model, also known as a *linguistic ontology*, was born within the Penman project [Matthiessen, 1987] as a fundamental resource for organizing domain knowledge appropriately for linguistic realization. It is a domain- and task-independent ontology meant to support and simplify the interface between domain knowledge and linguistic resources [Bateman, 1990]. The importance of this interface is clear if we think that most ideational inquiries ask questions regarding the classification of an input category in terms of abstract semantic categories. The Upper Model is based on the *Bloomington Lattice* [Matthiessen, 2005, page 168], an *ideational* grammatical semantic typology for English started by Michael Halliday and Christian Matthiessen during the summer of 1986. It reflects the English lexicogrammatical semantics, the *ideational* metafunction only, and it is called the *ideation base*. Figure 16 shows an excerpt of the higher level classes of the Penman Upper model and their taxonomical relations.

The Penman Upper Model was augmented to account for the grammar of German in the 1990's and it became the *Merged Upper Model* [Henschel, 1993; Henschel and Bateman, 1994] for use in the KOMET-Penman Multilingual Development Environment (KPML) system (see Section 6.5 below).

In order to provide more linguistic coverage, both in terms of the generation ability in a given language, but also in various other languages, and to bring the Merged Upper Model more in line with the systemic work of Halliday and Matthiessen [Halliday and Matthiessen, 1999], the *Generalized Upper Model* (GUM) [Bateman *et al.*, 1995] was created. At the time of writing, the latest

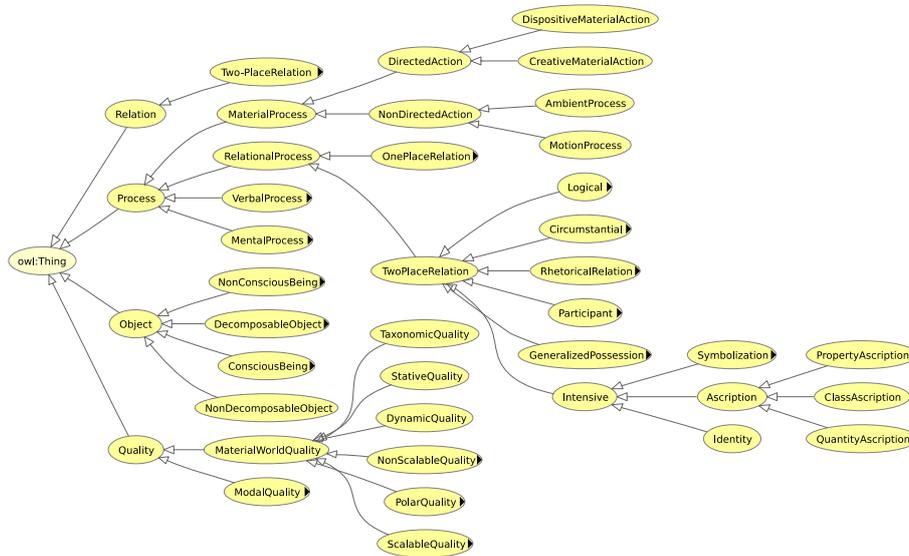


Figure 16: Excerpt of the Penman Upper Model taxonomy

version of GUM is 3.0<sup>6</sup>. In terms of representation format, the GUM, originally written in LOOM [MacGregor and Bates, 1987], has been made available as an OWL-DL<sup>7</sup> file.

The contents of the Penman Upper Model were used in conjunction with the Sentence Plan Language (SPL) [Kasper, 1989] (presented below) as input to the Penman generation system. The KPML system, instead, employs the Merged Upper Model (LOOM format), hereinafter referred to as the Upper Model..

#### 6.4 Input specification: the Sentence Plan Language

An SPL representation is a list of terms which describe the entities that need to be expressed in NL along with the particular attributes of those entities. Attributes may specify semantic relations that are to be expressed from the domain model or they may specify responses to inquiries about grammatical features of sentences. The syntax of SPL, specified in [Kasper, 1989], is reported here:

*Plan* → *Term*<sup>+</sup>  
*Term* → (*Variable* / *Type Attribute*<sup>\*</sup>) | *Variable* | *Constant* |  
(*Term*<sup>+</sup>) | (: *and Term*<sup>+</sup>) | (: *or Term*<sup>+</sup>)  
*Type* → *ConceptName* | (*ConceptName*<sup>+</sup>)  
*Attribute* → *Keyword Term*  
*Keyword* → *RelationName* | *MacroName* | *InquiryName* (*Variable*<sup>+</sup>) |  
*SpecialKeyword*

<sup>6</sup><http://pur1.org/net/gum>

<sup>7</sup><http://www.w3.org/TR/owl-guide>

**Example 2.** *A sentence plan for “The car is equipped with a service booklet.” could be:*

```
(e1 / be-equipped-with
 :actor (e2 / car
        :determiner the)
 :actee (e3 / service-booklet))
```

e1 represents the main term of this plan, and it denotes an entity of the domain model. The type of e1 is be-equipped-with which is defined as specialization of generalized-possession, a reified relation from the upper model. It has two main attributes named :actor and :actee. The actor is denoted by the variable e2 (referring to the concept car) and the actee by e3 (referring to the concept service-booklet).

The syntax of SPL permits the use of macros as keywords also. With a macro we can express in a succinct manner a set of delicate features to generate some specific grammatical phenomenon. E.g. if we want to express English tense in general terms, we should provide precise inquiry responses, setting three times and the ordering relations among them:

- the actual speaking time
- the event time, and
- the time of reference with which the event is contrasted.

The :tense macro was created to simply avoid specifying these temporal relations, simply distinguishing the English tenses using values that are expanded into the appropriate inquiry responses.

KPML contains a package of macro keywords that greatly help in simplifying the creation of a sentence plan.

**Example 3.** *To modify the previous sentence plan in order to generate the sentence “The car was equipped with a service booklet”, we can use the :tense macro as follows:*

```
(e1 / be-equipped-with
 :actor (e2 / car
        :determiner the)
 :actee (e3 / service-booklet)
 :tense past)
```

Another facility provided by KPML is a way of defining *default values* for inquiries in order to predefine sentence features that do not change frequently in a given application domain. The SP given above doesn't specify if the sentence to be generated must be a statement, a question, or a command, nor does it say if it should have positive or negative polarity. This is because the KPML system generates by default statements with positive polarity.

The interpretation of a SP is done by KPML in two phases:

1. the plan is first transformed into an internal representation, where all macros are expanded, and type information is distributed to variable terms whose consistency is also checked. The first term of the plan is treated as the initial unit to be expressed (main clause of the sentence).

2. given this representation, the generation process is guided by means of a series of inquiries to the sentence plan and the available knowledge sources according to this order:
  - (a) SPL keyword: The SP is searched first for a keyword matching the name of the inquiry, and the value is returned;
  - (b) knowledge sources: inquiries may have a function associated with them called the *inquiry implementation*, which searches the domain and upper model for the type or the attributes of the SPL terms;
  - (c) active default value: If an undefined answer is supplied by the inquiry implementation, or if there is no inquiry implementation, then the current active default value is used.

## 6.5 The KPML System

In this section, proceeding from the introduction of Section 5.2, we present the architecture of the KPML system along with an overview of the generation process (based on [Bateman, 1997b]) and the resources that the system uses for this purpose.

### 6.5.1 The KPML Generation Process

KPML uses a Penman-style generation architecture that is depicted in Figure 17. Generation in KPML proceeds in cycles of traversal through the system network. The outcome of this traversal is a set of grammatical features called *selection expressions* and a resulting grammatical structure. It is by resolving grammatical constraints associated with features of the selection expression that the grammatical structure is created. Features chosen during network traversal are selected according to the semantic input that needs to be expressed, an operation that is mediated by the chooser and inquiry framework (see Section 6.2): Choosers organize inquiries into “decision trees”, and inquiries are responsible for (a) inspecting the semantic specification that is being expressed in order to classify it and (b) providing access to particular portions of the semantic specification in order to trigger further realization. The connection between the systemic grammar and the semantic input is made via a *function association table* that relates grammatical functions (labels for grammatical constituents) and semantic “hubs” (labels for the semantic input chunks that need to be expressed). The input arguments for inquiries are grammatical functions.

Cycles of generation will continue for all sub-constituents of a grammatical unit until all sub-constituents are filled by some linguistic substance, usually *lexemes* or *morphemes*. One thing that has to be avoided is underconstraining grammatical constituents, which would cause infinite regression. There are four ways in KPML to correctly specify a grammatical constituent so that it receives lexical material and doesn’t trigger another cycle through the grammar:

1. an explicit lexical entry can be selected with the realization statement `lexify`;
2. a set of lexical features can be associated with a grammatical constituent using the `classify` realization statement; on completion of a traversal

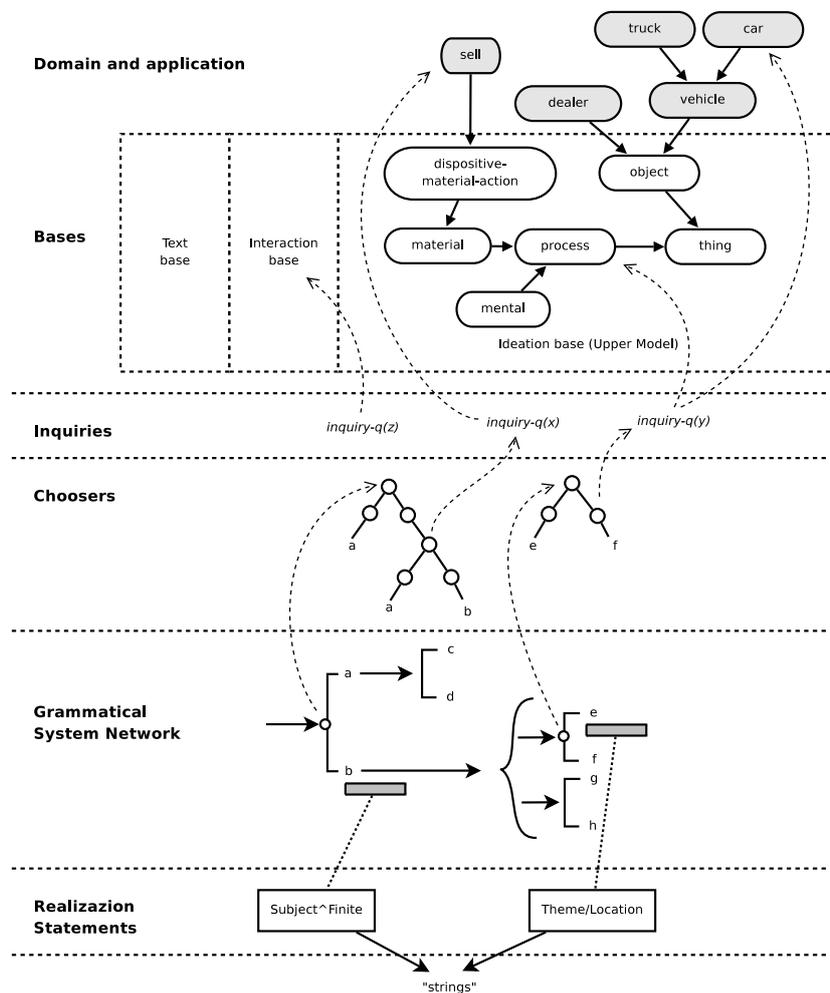


Figure 17: KPML Penman-style generation architecture (based on [Bateman, 1997b])

through the grammar, the complete collection of lexical features for a grammatical constituent is used to pick a matching lexical item (i.e. a lexical item whose features unify);

3. the inquiry `term-resolve-id` can be invoked to ask for an explicit lexicalization on semantic grounds;
4. an explicit selection of a morpheme can be made with the morphological realization operators, which are: `preselect-substance`, `preselect-substance-as-stem`, or `preselect-substance-as-property`.

It must be noted that if a constituent has been classified, the selection of a lexical item as shown in (2) will not respect any additional information since it follows a purely lexicogrammar internal selection. This means that no semantic information or SPL input will be consulted. If we need to take into

account semantic information also, option (3) must be chosen by including the `term-resolve-id` inquiry in some chooser that is activated at an appropriate point during generation.

The semantic organization adopted by `KPML` foresees first of all a linguistic ontology called the *Upper Model* that was presented in Section 6.3. All of the `KPML` resources are defined in a way that generation is possible with respect to a single Upper Model, as concrete instantiation of the *ideation base*. The domain model representing the universe of discourse we want to generate natural language about, must be connected with the upper model. This way we can directly use entities from the domain model to formulate SPL inputs for the generator. Two other “bases” are needed (as shown in Figure 17): Interaction and text base. The *interaction base* represents the knowledge that the system has about the social and epistemic relationship between speaker (machine) and the hearer; this can be instantiated as a *user model*. The *text base* instead, is concerned with the system’s knowledge about which discourse structures, coherence relations, and cohesive ties need to be used, which grammatically are interpreted as theme-rheme structure, conjunctions, referring-expressions, etc.

### 6.5.2 KPML Input Resources

In order to be able to generate, the `KPML` system needs the following linguistic resources:

- a domain model,
- a grammar,
- and a lexicon.

We introduce them briefly hereinafter, suggesting the reader to refer to [Bateman, 1997b, Section 12.2] for an in-depth description of resource organization and definition formats.

**Domain Model** Given a domain model on which we want to generate, its concepts and properties (relations and attributes) must be subordinated to the Upper Model entities by means of LOOM axioms. This means we have to rewrite the original domain ontology in the input format required by `KPML` (LOOM), subordinating it to the Upper Model. It has to be noted that `KPML` is now moving towards OWL-DL [W3C, 2008] and will soon untie its dependency on LOOM. For generation purposes, not all axioms of the original ontology need to be translated, but just concept and role (abstract and concrete) definitions, and subsumption relations. The mapping is quite simple, since all source entities (both concepts and roles) will be translated into LOOM concepts, and either subordinated to an `UM Object` or a `Process` or one of their descendants (see Table 8).

Since attribute descriptions will be rendered using a copula (e.g. “the engine’s power *is* 250 HP”, “the car’s weight *is* 1500 kg”), an `UM Class Ascription` (`Process`) needs to be used in the sentence plan to describe this process. Furthermore the reified attribute (see Table 8) will most probably become the subject of

Domain Ontology Entity	... mapped into a	... subordinated by an UM
concept	concept	Object
relation <sub>v</sub>	concept	Process
relation <sub>a</sub>	concept	Object
attribute	concept	Object

Table 8: Mapping of domain ontology entities and subordination to UM entities

the clause, and the domain concept of the attribute will be used as the subject's modifier.

Relations instead, have been distinguished into two categories: relation<sub>v</sub> and relation<sub>a</sub>. The former represents relations which will be expressed as verbs describing the respective processes as in "the car *runs* on gasoline". The latter refers to those relations that act as attributes (but have a concept as range instead of a concrete datatype) and are treated the same way as attributes are (e.g. "the car's *make* is VW", "the car's *model* is Golf GTD").

```
(defconcept Vehicle
  :is (:and Penman-kb::Decomposable-Object :primitive))
(kpml::annotate-concept Vehicle :lex-items (vehicle))

(defconcept Car
  :is (:and Vehicle :primitive))
(kpml::annotate-concept Car :lex-items (car))

(defconcept Car-Dealer
  :is (:and Penman-kb::Object :primitive))
(kpml::annotate-concept Car-Dealer :lex-items (car-dealer))

(defconcept Make
  :is (:and Penman-kb::Object :primitive))
(kpml::annotate-concept Car :lex-items (make))

(defconcept Model
  :is (:and Penman-kb::Object :primitive))
(kpml::annotate-concept Model :lex-items (model))

(defconcept Sell
  :is (:and Penman-kb::Dispositive-Material-Action :primitive))
(kpml::annotate-concept Sell :lex-items (sell))
```

Figure 18: Excerpt of LOOM ontology from the automotive domain

Some example concept definitions in LOOM format regarding the automotive domain are reported in Figure 18. The `:primitive` predicate means that the concept it refers to is incompletely specified, i.e. there are hidden attributes

about objects of that type that are not represented and the concept is thus considered as a ‘primitive’. The classifier will not try and put it somewhere else on the basis of any features it may or may not have. This is necessary because otherwise it would be unified with other concepts from which it is not formally differentiated. *v*-type relations of the original domain ontology must be reified and represented as concepts subsumed by one of the “processes” of the Upper Model (see Figure 16). The remaining *a*-type relations are reified and classified under the Object sub-hierarchy (see concepts *Make* or *Model* above) and, as stated above, will be usually rendered as subject of the generated clause (e.g. “the car’s *make* is Opel”). The domain concept of the relation will be used as subject modifier (as *car’s* in the previous example).

The `kpml::annotate-concept` lines following concept definitions are necessary to create a link between the defined concept and the respective lexical items contained in the Lexicon.

**Grammar** Although in our work we currently generate in English using the latest Nigel Grammar for English (Section 6.2), it is good to know that there are other resources available for a range of languages (including resources for Chinese, Czech, Greek, Japanese, Russian, German, and Spanish in varying stages of development)<sup>8</sup>. The Nigel Grammar for English consists of 42 functional regions (see Table 9), each one giving its name to three files, one containing the systems, one the choosers and the third one the enquiries for that region. E.g. the RANKING region is covered by the three files which are RANKING.systems, RANKING.choosers, and RANKING.inquiries.

ADJECTIVAL-COMPARISON	EPITHET	PROCESSUALTHINGTYPE
ADJECTIVAL-GROUP	GATES	PRONOUN
ADVERBIAL	MOOD	QUALIFICATION
ATTITUDE	NOMINALGROUPCOMPLEXITY	QUANTIFICATION
CIRCUMSTANTIAL	NOMINAL-PERSON	QUANTITY-GROUP
CLASSIFICATION	NONRELATIONALTRANSIT.	RANKING
CLAUSECOMPLEX	NOUNTYPE	RELATIONALTRANSITIVITY
CONJUNCTION	ORDINALITY	SELECTION
COUNTNUMBER	PHRASAL-MOOD	TAG
CULMINATION	POLARITY	TENSE
DEPENDENCY	POST-DEICTICITY	THEME
DETERMINATION	PPCOMPLEXITY	UNIFYINGGATES
ELABORATION	PPOTHER	VOICE
ELLIPSIS	PPSPATIOTEMPORAL	WORD-FORMS

Table 9: Functional regions in the Nigel Grammar for English

<sup>8</sup>A collection of systemic-functional grammars for natural language generation can be found in the Generation Bank of the University of Bremen. The generation bank is a website that is being constructed to contain lexicogrammars for tactical generation in a variety of languages. All grammars have the same form and can be used by the same generator i.e. the KPML system. When complete, each grammar fragment will contain a complete grammar definition in KPML-standard format, including example sets (‘target suites’) that provide a summary of coverage and corresponding semantic inputs. It’s available here: <http://www.fb10.uni-bremen.de/anglistik/langpro/kpml/genbank/generation-bank.html>.

The grammar files are written using LISP-like syntax, as shown in the following triplet of system, chooser, and inquiry taken from the MOOD region.

```
(SYSTEM
  :NAME      MOOD-TYPE
  :INPUTS    INDEPENDENT-CLAUSE-SIMPLEX
  :OUTPUTS
    ((0.5 INDICATIVE)
     (0.5 IMPERATIVE
      (INSERT NONFINITIVE)
      (INFLECTIFY NONFINITIVE STEM)))
  :CHOOSER   MOOD-TYPE-CHOOSER
  :REGION    MOOD
  :METAFUNCTION  INTERPERSONAL
)
```

```
(CHOOSER
  :NAME      MOOD-TYPE-CHOOSER
  :DEFINITION
    ((ASK (COMMAND-Q SPEECHACT)
          (COMMAND
           (IDENTIFY SUBJECT
            (COMMAND-RESPONSIBLE-ID SPEECHACT))
           (CHOOSE IMPERATIVE))
          (NOCOMMAND
           (CHOOSE INDICATIVE))))
)
```

```
(ASKOPERATOR
  :NAME      COMMAND-Q
  :DOMAIN    TP
  :PARAMETERS (ACT1)
  :ENGLISH   (
    "  Is the illocutionary point of the surface level"
    "speech act represented by"
    ACT1
    "  a command, i.e. a request of an action by the"
    "hearer?"
  )
  :OPERATORCODE  KPML::COMMAND-Q-CODE
  :PARAMETERASSOCIATIONTYPES (CONCEPT)
  :ANSWERSET     (COMMAND NOCOMMAND)
)
```

Finally we see how lexical items are stored and described in KPML.

**Lexicon** Three lexical items taken from the automotive domain are presented in Figure 19.

```

(lexical-item
  :name car
  :spelling "car"
  :sample-sentence "This car is a Land Rover."

  :features (common-noun noun)
  :editor "PAOLO DONGILLI"
  :date "Mon Sep 19 15:26:39 CEST 2006"
)

(lexical-item
  :name sell
  :spelling "sell"
  :sample-sentence "This car is sold by a car dealer."
  :features (disposal-verb do-verb effective-verb s-irr)
  :properties ((pastform "sold")
               (edparticipleform "sold"))
  :editor "PAOLO DONGILLI"
  :date "Wed Sep 14 11:11:39 CEST 2006"
)

(lexical-item
  :name car-dealer
  :spelling "car dealer"
  :sample-sentence "The car dealer is located in Germany."
  :features (common-noun noun)
  :editor "PAOLO DONGILLI"
  :date "Wed Sep 14 10:59:45 CEST 2006"
)

```

Figure 19: Lexical items from the automotive domain

The features that appear under the features slot depend on the concrete linguistic resources defined to the system. The `properties` slot, instead, is used for holding idiosyncratic exceptions to general morphological processes. A resource-external morphology handling is adopted in `KPML`, i.e. the resource definitions assume that the morphological features that they use are interpreted by some non-systemic component of `KPML`. One example of such a resource definition is the Nigel grammar of English, for which the Penman system provided hardcoded English morphology. This hardcoded morphology is inherited by `KPML`<sup>9</sup>.

<sup>9</sup>The current version of the Nigel grammar released as a `KPML`-resource set, does however include systemic resources for morphology. This provides a more flexible and transparent representation of the linguistic resources at word and morpheme rank, but increases the generation time a little since further cycles through the grammar are required.

## 7 Conclusions

This work presented our efforts in terms of devising an NLG pipelined architecture to render in natural language a given conjunctive query over a Description Logics knowledge base. We analyzed in depth the main modules of our architecture, including, if seen at high level, *text planning*, *sentence planning*, and *linguistic realization* (otherwise called *surface realization*). Each one of these high-level phases was further split into its subtasks.

We started presenting six possible strategies for *discourse planning* of a given complex concept description, concentrating on three different goals:

1. maximization of local referential-coherence (CT);
2. minimization of overall conceptual distance (mCD);
3. minimization of change in the discourse plan between consecutive edits (user-driven, depth-first, relation-priority depth-first).

After showing the pros and cons of every approach, we chose the relation-priority depth-first strategy, recognizing in the usage of relation priorities a possible encoding of *domain communication knowledge*.

Under *sentence planning* we focused first on *sentence aggregation* proposing several aggregation templates and Algorithm 8 for calculating the best covering match. As far as *referring expression generation* is concerned, after constraining the number of aggregation expressions needed, we proposed Algorithm 9 for the assignment of appropriate referring expressions for each entity present in a given text plan.

With the outputs obtained from the discourse planning, sentence aggregation, and referring expression generation phases, we described in detail how to employ these pieces of information to correctly generate the input for the linguistic realizer. The input is called *sentence plan* and the language used is the *Sentence Plan Language* (SPL).

At last we described the *linguistic realization* phase along with its possible approaches. Among the feature-based systems which guarantee the highest level of sophistication and flexibility, we chose to employ and describe the KPML system, a free multilingual systemic-functional grammar (SFG) development application and linguistic realizer that uses SPL as input language.

With a whole NLG pipeline in place we are now able to start implementing a new interface (described in [Dongilli *et al.*, 2006]) for our intelligent query tool, an interface with natural language feedback according to the WYSIWYM paradigm [Power and Scott, 1998].

# Appendix

## English semantics for generation with KPML

We report here for commodity the introductory guide by Juan Rafael Zamorano Mansilla [Zamorano Mansilla, 2003] which provides details about the possible forms that can be used when building SPL specifications for driving natural language generation with KPML.

The main kinds of information that are necessary when generating sentences from SPL specifications are the following:

- Choice of **process type**: this determines the basic semantic domain of the specification;
- Choice of **circumstances**: these provide additional information concerning when, where, how, etc. the process took place;
- Choice of **sentence types**: this specifies forms of sentences according to their intended function in discourse.

### Process types

The following tree shows the process types defined in the Merged Upper Model and usable in SPL input specifications for generating natural language with the KPML system. The most common process types are shown in boldface, and followed by the participants (in brackets) that are inherent to them. An example sentence illustrates an instance of that process type.

#### 1. RELATIONAL-PROCESS

##### 1.1. ONE-PLACE-RELATION

1.1.1. **EXISTENCE** (:domain) *There is a book on the table.*

##### 1.2. TWO-PLACE-RELATION

1.2.1. **GENERALIZED-POSSESSION** (:domain, :range) *I've got two brothers.*

1.2.2. **CIRCUMSTANTIAL** (:domain, :range) *The stone weighs eight kilos.*

##### 1.2.3. INTENSIVE

###### 1.2.3.1. ASCRIPTION

1.2.3.1.1. **PROPERTY-ASCRIPTION** (:domain, :range) *My tailor is rich.*

1.2.3.1.2. **CLASS-ASCRIPTION** (:domain, :range) *My father is a teacher.*

1.2.3.1.3. QUANTITY-ASCRIPTION

###### 1.2.3.2. UM-IDENTITY

###### 1.2.3.3. SYMBOLIZATION

#### 2. MENTAL-PROCESS

##### 2.1. MENTAL-ACTIVE

## 2.2. MENTAL-INACTIVE

2.2.1. **COGNITION** (:senser, :phenomenon) *I know the answer.*

2.2.1.1. BELIEVE

2.2.1.2. KNOW

2.2.1.3. THINK

2.2.2. **REACTION** (:senser, :phenomenon) *I don't like tea.*

2.2.2.1. LIKING

2.2.2.2. STRIVING

2.2.2.3. WANTING

2.2.2.4. DISLIKING

2.2.2.5. FEARING

2.2.3. **PERCEPTION** (:senser, :phenomenon) *Nobody saw the accident.*

## 3. VERBAL-PROCESS

3.1. **ADDRESSEE-ORIENTED-VERBAL-PROCESS** (:sayer, :saying, :recipient) *I told her the news.*

3.2. **NON-ADDRESSEE-ORIENTED-VERBAL-PROCESS** (:sayer, :saying) *I didn't say that.*

## 4. MATERIAL-PROCESS

### 4.1. DIRECTED-ACTION

4.1.1. **CREATIVE-MATERIAL-ACTION** (:actor, :actee, :beneficiary) *My brother has written a book.*

4.1.2. **DISPOSITIVE-MATERIAL-ACTION** (:actor, :actee, :beneficiary) *We have changed the first chapter.*

4.2. **NONDIRECTED-ACTION** (:actor) *He died.*

4.2.1. **AMBIENT-PROCESS** (usually no participant involved) *It's raining.*

## Circumstances

This is a list of the circumstances recognized by the semantic organization built into KPML. Most circumstances types are defined just like participants: first you type the circumstance type after a colon (shown in blue below), and then in brackets you write the name of the circumstance, the semantics of the constituent that comes with the preposition and the rest of information (which is the same as for participants, because you always have participants or processes after prepositions). Other circumstances however have a more complex formalism, including two names and two places for semantics. NAME1 refers to the name given to the circumstance relation, while NAME2 refers to the name of the participant that comes with the preposition. Logically, the semantics relative to that participant are placed next to NAME2. Sometimes, however, it is important to specify the precise semantics to obtain the right generation. In these cases the semantics appears in bold. Examples, again drawing on the results that would be produced by the Nigel grammar of English, are:

```

:inclusive (accom-1 / object :lex money)
:destination (Fr / object :lex France)
:absolute-temporal-extent (tempo / object :lex day :number plural)
:matter-q matter
:matter-id
  (ab01 / empty :domain x :range (book / object :lex book :determiner the))

```

In red you can find the result of generation with these commands.

#### ACCOMPANIMENT

- **with**  
:inclusive ([name] / [semantics] :lex [item])
- **as well as**  
:additive ([name] / [semantics] :lex [item])
- **instead of**  
alternative ([name] / [semantics] :lex [item])
- **without**  
:exclusive ([name] / [semantics] :lex [item])

#### CAUSE

- **because of**  
:reason ([name] / [semantics] :lex [item])
- **for (purpose)**  
:purpose ([name] / [semantics] :lex [item])
- **for (client)**  
:client ([name] / [semantics] :lex [item])
- **in spite of**  
:causal-relation ([name] / [semantics] :lex [item])

#### COMPARISON

- **like**  
:similarity ([name] / [semantics] :lex [item])
- **similar to**  
:know-manner-q known  
:process-manner-id ([name1] / [semantics]  
:resemblance-q resemblance  
:formal-register-q formal  
:concrete-comparison-q concrete  
:domain x  
:range ([name2] / [semantics of the participant that comes with  
the preposition] ))
- **different from**  
:know-manner-q known  
:process-manner-id ([name1] / [semantics]  
:resemblance-q resemblance  
:resemblance-type-q difference  
:domain x  
:range ([name2] / [semantics of the participant that comes with  
the preposition] ))

## MEANS

- **Adverbial Group**  
:manner ([name] / [semantics] :lex [item])
- **by (generalized means)**  
:generalized-means ([name] / [semantics] :lex [item])
- **by (enablement)**  
:enablement ([name] / [semantics] :lex [item])
- **by (agentive)**  
:agentive ([name] / [semantics] :lex [item])
- **by means of**  
:know-manner-q known  
:process-manner-id ([name1] / enablement  
:explicit-means-q explicit  
:domain x  
:range ([name2] / [semantics of the participant that comes with  
the preposition] ))
- **with (instrumental)**  
:instrumental ([name] / [semantics] :lex [item])

## SUBJECT-MATTER

- **concerning**  
:specific-matter ([name] / [semantics] :lex [item])
- **in the case of**  
:matter-q matter  
:matter-id ([name1] / specific-matter  
:matter-coverage-q clause  
:domain x  
:range ([name2] / [semantics of the participant that comes with  
the preposition] ))
- **about** :diffuse-matter ([name] / [semantics] :lex [item])
- **as to**  
:matter-q matter  
:matter-id ([name1] / diffuse-matter  
:matter-coverage-q clause  
:domain x  
:range ([name2] / [semantics of the participant that comes with  
the preposition] ))
- **of**  
:matter-q matter  
:matter-id ([name1] / diffuse-matter  
:formal-register-q formal  
:domain x  
:range ([name2] / [semantics of the participant that comes with  
the preposition] ))

## ROLE-PLAYING

- as :role-playing ([name] / [semantics] :lex [item])

#### TEMPORAL EXTENT

- **for (temporal extent)**  
:absolute-temporal-extent ([name] / [semantics] :lex [item])
- **in (temporal extent)**  
:relative-temporal-extent ([name] / [semantics] :lex [item])
- **during**  
:exhaustive-duration ([name] / [semantics] :lex [item])

#### SPATIAL EXTENT

- **for (spatial extent)**  
:absolute-spatial-extent ([name] / [semantics] :lex [item])
- **along**  
:parallel-extent ([name] / [semantics] :lex [item])
- **across**  
:nonparallel-extent ([name] / [semantics] :lex [item])

#### SPATIAL LOCATION

- **Adverbial Group**  
:spatial-location-specification-q spatiallocation  
:spatial-location-id ([name1] / [semantics]  
:identifiability-q identifiable  
:location-relation-specificity-q unspecified  
:lex [item])
- **at (spatial location)**  
:spatial-locating ([name] / space-point :lex [item])
- **in (spatial location)**  
:spatial-locating ([name] / three-d-location :lex [item])
- **outside**  
:spatial-location-specification-q spatiallocation  
:spatial-location-id ([name1] / [semantics]  
:containment-q noncontainment  
:domain x  
:range ([name2] / three-d-location :lex [item] ))
- **inside**  
:spatial-location-specification-q spatiallocation  
:spatial-location-id ([name1] / [semantics]  
:explicit-containment-q explicit  
:domain x  
:range ([name2] / three-d-location :lex [item] ))
- **on**  
:spatial-locating ([name] / one-or-two-d-location :lex [item])
- **beside**  
:horizontal ([name] / [semantics] :lex [item])

- **next to**
  - :spatial-location-specification-q spatiallocation
  - :spatial-location-id ([name1] / horizontal
    - :immediate-adjacency-q adjacent
    - :specify-adjacency-q specified
    - :domain x
  - :range ([name2] / [semantics] :lex [item] ))
- **between**
  - :between ([name] / [semantics] :lex [item])
- **behind**
  - :behind ([name] / [semantics] :lex [item])
- **below**
  - :below ([name] / [semantics] :lex [item])
- **underneath**
  - :spatial-location-specification-q spatiallocation
  - :spatial-location-id ([name1] / below
    - :area-of-coverage-q partial
    - :domain x
  - :range ([name2] / [semantics] :lex [item] ))
- **under**
  - :spatial-location-specification-q spatiallocation
  - :spatial-location-id ([name1] / below
    - :area-of-coverage-q partial
    - :surface-contact-q noncontact
    - :domain x
  - :range ([name2] / [semantics] :lex [item] ))
- **above**
  - :above ([name] / [semantics] :lex [item])
- **over**
  - :spatial-location-specification-q spatiallocation
  - :spatial-location-id ([name1] / above
    - :area-of-coverage-q partial
    - :surface-contact-q noncontact
    - :domain x
  - :range ([name2] / [semantics] :lex [item] ))
- **on top of**
  - :spatial-location-specification-q spatiallocation
  - :spatial-location-id ([name1] / above
    - :area-of-coverage-q partial
    - :domain x
  - :range ([name2] / [semantics] :lex [item] ))
- **in front of**
  - :facing ([name] / [semantics] :lex [item]) to (destination)
  - :destination ([name] / [semantics] :lex [item])
- **onto**
  - :spatial-location-specification-q spatiallocation
  - :spatial-location-id ([name1] / destination

- :domain x
  - :range ([name2] / one-or-two-d-location :lex [item] ))
- **into**
  - :spatial-location-specification-q spatiallocation
  - :spatial-location-id ([name1] / destination
  - :domain x
  - :range ([name2] / three-d-location :lex [item] ))
- **towards**
  - :spatial-location-specification-q spatiallocation
  - :spatial-location-id ([name1] / destination
  - :orientation-q oriented
  - :domain x
  - :range ([name2] / [semantics] :lex [item] ))
- **from**
  - :source ([name] / [semantics] :lex [item])
- **off**
  - :spatial-location-specification-q spatiallocation
  - :spatial-location-id ([name1] / source
  - :domain x
  - :range ([name2] / one-or-two-d-location :lex [item] ))
- **out of**
  - :spatial-location-specification-q spatiallocation
  - :spatial-location-id ([name1] / source
  - :domain x
  - :range ([name2] / three-d-location :lex [item] ))
- **away from**
  - :spatial-location-specification-q spatiallocation
  - :spatial-location-id ([name1] / source
  - :orientation-q oriented
  - :domain x
  - :range ([name2] / [semantics] :lex [item] ))

## TEMPORAL LOCATION

- **Adverbial Group**
  - :temporal-location-specification-q temporallocation
  - :temporal-location-id ([name1] / [semantics]
  - :identifiability-q identifiable
  - :location-relation-specificity-q unspecified
  - :lex [item])
- **at**
  - :temporal-locating ([name] / time-point :lex [item])
- **in**
  - :temporal-locating ([name] / three-d-time :lex [item])
- **on**
  - :temporal-locating ([name] / one-or-two-d-time :lex [item])
- **by (temporal location)**
  - :temporal-ordering ([name] / [semantics] :lex [item])

- **before**  
:anterior ([name] / [semantics] :lex [item])
- **until**  
:temporal-location-specification-q temporallocation  
:temporal-location-id ([name1] / anterior  
:period-extremal-q periodextremal  
:domain x :range ([name2] / [semantics] :lex [item] ))
- **after**  
:posterior ([name] / [semantics] :lex [item])
- **since**  
:temporal-location-specification-q temporallocation  
:temporal-location-id ([name1] / posterior  
:period-extremal-q periodextremal :domain x :range  
([name2] / [semantics] :lex [item] ))
- **from**  
:temporal-location-specification-q temporallocation  
:temporal-location-id ([name1] / posterior  
:period-extremal-q periodextremal  
:period-time-or-state-q state-or-activity  
:domain x  
:range ([name2] / [semantics] :lex [item] ))

### Sentence Types

Many of these specifications are actually 'macros', which means that they are shorthand for something more complex. They allow to write simple SPLs without worrying about the semantic specification that is behind them. For example, ":tense present" produces a sentence in the simple present tense without requiring that you know that this shorthand for a specific set of temporal relations between the time of speaking and the time of the event described. There are occasions when you need to delve more deeply, but for a beginning you can often get by without.

- **WH-QUESTIONS**  
:speech-act-id (q / question :polarity positive)  
:question-item-id [name of the participant or circumstance asked about]
- **YES/NO-QUESTIONS**  
:speech-act-id (q / question :polarity variable)
- **TENSE**  
:tense [present, past, future, present-continuous (or present-progressive),  
past-continuous (or past-progressive), future-continuous (or future-progressive),  
present-perfect, past-perfect, future-perfect, future-in-present (going to),  
present-perfect-continuous]
- **VOICE**  
:voice [active/passive]
- **POLARITY**  
:polarity [positive/negative]

- **MODALITY**

:modality [can, cant, could, couldnt, may, might, must, neednt, should, shouldnt, will, wont, would, wouldnt]

## References

- [Baader *et al.*, 2007] Franz Baader, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, 2nd edition, 2007.
- [Bateman *et al.*, 1995] John A. Bateman, Bernardo Magnini, and Giovanni Fabris. The generalized upper model knowledge base: Organization and use. In N. J. I. Mars, editor, *Towards very large knowledge bases: knowledge building and knowledge sharing*, pages 60–72, Amsterdam, 1995. IOS Press.
- [Bateman, 1990] John A. Bateman. Upper modeling: organizing knowledge for natural language processing. In *Proceedings of the Fifth International Natural Language Generation Workshop*, pages 54–60, Pittsburgh, PA., 1990. Organized by Kathleen R. McKeown (Columbia University), Johanna D. Moore (University of Pittsburgh) and Sergei Nirenburg (Carnegie Mellon University). Held 3-6 June 1990, Dawson, PA.
- [Bateman, 1997a] John A. Bateman. Enabling technology for multilingual natural language generation: the KPML development environment. *Journal of Natural Language Engineering*, 3(1):15–55, 1997.
- [Bateman, 1997b] John A. Bateman. *KPML Development Environment: multilingual linguistic resource development and sentence generation*. German National Center for Information Technology (GMD), Institute for integrated publication and information systems (IPSI), Darmstadt, Germany, January 1997. (Release 1.1).
- [Brennan *et al.*, 1987] Susan E. Brennan, Marilyn W. Friedman, and Carl J. Pollard. A centering approach to pronouns. In *Proceedings of the 25th annual meeting on Association for Computational Linguistics*, pages 155–162, Morristown, NJ, USA, 1987. Association for Computational Linguistics.
- [Cheng *et al.*, 1997] Hua Cheng, Chris Mellish, and Michael O’Donnell. Aggregation based on text structure for descriptive text generation. In *Proceedings of the PhD Workshop on Natural Language Generation, 9th European Summer School in Logic, Language and Information (ESSLLI97)*, Aix-en-Provence, France, 1997.
- [Dongilli *et al.*, 2004] Paolo Dongilli, Enrico Franconi, and Sergio Tessaris. Semantics Driven Support for Query Formulation. In *Proceedings of the 2004 International Workshop on Description Logics (DL 2004)*, volume 104, Whistler, BC, Canada, June 2004.
- [Dongilli *et al.*, 2006] Paolo Dongilli, Sergio Tessaris, and John Bateman. Leveraging Systemic-Functional Linguistics to Enhance Intelligent Database Querying. In *Proceedings of the Sixth International Conference on Intelligent Systems Design and Applications*, Jinan, China, October 2006.

- [Dongilli, 2007a] Paolo Dongilli. Centering-theory-based text planning of a conjunctive query. In *Proceedings of the 3rd Language & Technology Conference*, Lecture Notes in AI (LNAI), Poznan, Poland, October 2007. Springer Verlag.
- [Dongilli, 2007b] Paolo Dongilli. Discourse planning strategies for complex concept descriptions. In *Proceedings of the 7th International Symposium on Natural Language Processing (SNLP-2007)*, Pattaya, Chonburi, Thailand, December 2007.
- [Galanis and Androutsopoulos, 2007] Dimitrios Galanis and Ion Androutsopoulos. Generating Multilingual Descriptions from Linguistically Annotated OWL Ontologies: the NaturalOWL System. In *Proceedings of the 11th European Workshop on Natural Language Generation (ENLG 2007)*, pages 143–146, Schloss Dagstuhl, Germany, 2007.
- [Grosz and Sidner, 1986] Barbara J. Grosz and Candace L. Sidner. Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3):175–204, 1986.
- [Grosz *et al.*, 1986] Barbara J. Grosz, Aravind K. Joshi, and Scott Weinstein. Towards a computational theory of discourse interpretation. Unpublished ms., 1986.
- [Grosz *et al.*, 1995] Barbara J. Grosz, Aravind K. Joshi, and Scott Weinstein. Centering: A framework for modeling the local coherence of discourse. *Computational Linguistics*, 21(2):203–225, 1995.
- [Halliday and Matthiessen, 1999] Michael A. K. Halliday and Christian M. I. M. Matthiessen. *Construing experience through meaning: a language-based approach to cognition*. Cassell, London, 1999.
- [Halliday and Matthiessen, 2004] Michael A. K. Halliday and Christian M.I.M. Matthiessen. *An Introduction to Functional Grammar*. Edward Arnold, London, 3rd edition, 2004.
- [Henschel and Bateman, 1994] Renate Henschel and John A. Bateman. The merged upper model: a linguistic ontology for German and English. In *Proceedings of COLING '94*, volume II, pages 803–809, Kyoto, Japan, August 1994.
- [Henschel, 1993] Renate Henschel. Merging the English and the German Upper Model. Technical report, GMD/Institut für Integrierte Publikations- und Informationssysteme, Darmstadt, Germany, 1993. Appears as: *Arbeitspapiere der GMD*, 848, June 1994. GMD, Sankt Augustin.
- [Horrocks and Tessaris, 2002] Ian Horrocks and Sergio Tessaris. Querying the semantic web: a formal approach. In Ian Horrocks and James Hendler, editors, *Proc. of the 2002 International Semantic Web Conference (ISWC 2002)*, number 2342 in Lecture Notes in Computer Science, pages 177–191. Springer-Verlag, 2002.
- [Hovy, 1997] Eduard Hovy. *Survey of the state of the art in human language technology*, chapter 4.1, Language generation: overview, pages 139–146. Cambridge University Press, New York, NY, USA, 1997.

- [Kasper, 1989] Robert T. Kasper. A flexible interface for linking applications to PENMAN's sentence generator. In *Proceedings of the DARPA Workshop on Speech and Natural Language*, 1989.
- [Kay, 1979] Martin Kay. Functional grammar. In Berkeley Linguistics Society, editor, *Proceedings of the 5th Meeting of the Berkeley Linguistics Society*, 1979.
- [Kibble and Power, 2004] Rodger Kibble and Richard Power. Optimizing referential coherence in text generation. *Computational Linguistics*, 30(4):401–416, 2004.
- [Kittredge *et al.*, 1991] Richard Kittredge, Tanya Korelsky, and Owen Rambow. On the need for domain communication knowledge. *Computational Intelligence*, 7(4):305–314, 1991.
- [MacGregor and Bates, 1987] Robert MacGregor and Raymond Bates. The LOOM knowledge representation language. In *Proceedings of the Knowledge-Based Systems Workshop, 1987*. Held in St. Louis, Missouri, April 21-23, 1987. Also available as ISI reprint series report, RS-87-188, USC/Information Sciences Institute, Marina del Rey, CA.
- [Mann and Matthiessen, 1983] William C. Mann and C. M. I. M. Matthiessen. Nigel: A systemic grammar for text generation. Technical Report RR-83-105, USC/Information Sciences Institute, February 1983. This paper also appears in a volume of the *Advances in Discourse Processes Series*, R. Freedle (ed.): *Systemic Perspectives on Discourse: Volume I*. published by Ablex.
- [Mann, 1983a] William C. Mann. The anatomy of a systemic choice. *Discourse Processes*, 1983. Also available as USC/Information Sciences Institute, Research Report ISI/RR-82-104, 1982.
- [Mann, 1983b] William C. Mann. An overview of the penman text generation system. In *AAAI*, pages 261–265, 1983.
- [Matthiessen, 1981] C. M. I. M. Matthiessen. A grammar and a lexicon for a text-production system. In *The Nineteenth Annual Meeting of the Association for Computational Linguistics*. Sperry Univac, 1981.
- [Matthiessen, 1983] Christian M. I. M. Matthiessen. The systemic framework in text generation: Nigel. In W. Greaves & J. Benson, editor, *Systemic Perspectives on Discourse*. Ablex, 1983.
- [Matthiessen, 1987] Christian M. I. M. Matthiessen. Notes on the organization of the environment of a text generation grammar. In Gerard Kempen, editor, *Natural Language Generation: Recent Advances in Artificial Intelligence, Psychology, and Linguistics*. Kluwer Academic Publishers, Boston/Dordrecht, 1987. Paper presented at the Third International Workshop on Natural Language Generation, August 1986, Nijmegen, The Netherlands.
- [Matthiessen, 2005] Christian M. I. M. Matthiessen. Remembering Bill Mann. *Computational Linguistics*, 31(2):161–172, 2005.
- [McRoy *et al.*, 2001] Susan W. McRoy, Songsak Channarukul, and Syed S. Ali. Creating natural language output for real-time applications. *Intelligence*, 12(2):21–34, 2001.

- [Melengoglou, 2002] Alexander Melengoglou. Multilingual aggregation in the m-piro system. Master's thesis, University of Edinburgh, 2002.
- [Mellish, 1988] Christopher S. Mellish. Implementing systemic classification by unification. *Journal of Computational Linguistics*, 14(1):40–51, 1988.
- [Merriam-Webster, 2007] Merriam-Webster. Merriam-Webster's Online Dictionary. <http://www.merriam-webster.com>, 2007.
- [Power and Scott, 1998] Richard Power and Donia Scott. Multilingual Authoring Using Feedback Texts. In *Proceedings of the 17th International Conference on Computational Linguistics and 36th Annual Meeting of the Association for Computational Linguistics (COLING-ACL 98)*, pages 1053–1059, Morristown, NJ, USA, 1998. Association for Computational Linguistics.
- [Reape and Mellish, 1999] Mike Reape and Chris Mellish. Just what *is* aggregation anyway? In *Proceedings of the 7th European Workshop on Natural Language Generation*, pages 20–29, Toulouse, 1999.
- [Reiter and Dale, 2000] Ehud Reiter and Robert Dale. *Building Natural Language Generation Systems*. Cambridge University Press, 2000.
- [Teich, 1999] Elke Teich. *Systemic Functional Grammar in Natural Language Generation: Linguistic Description and Computational Representation*. Cassell, London, 1999.
- [W3C, 2008] W3C. OWL Web Ontology Language Overview. <http://www.w3.org/TR/owl-features>, 2008. W3C Recommendation 10 February 2004.
- [Wilkinson, 1995] John Wilkinson. Aggregation in natural language generation: Another look. Co-op work term report, September 1995.
- [Zamorano Mansilla, 2003] Juan Rafael Zamorano Mansilla. Building SPL logical specifications for generation with kpml. <http://www.fb10.uni-bremen.de/anglistik/langpro/kpml/Doc/tutorial/enseman.html>, 2003.
- [Zorzi *et al.*, 2007] Ivan Zorzi, Sergio Tessaris, and Paolo Dongilli. Improving Responsiveness of Ontology-Based Query Formulation. In *Proceedings of the 4th Italian Workshop on Semantic Web Applications and Perspectives (SWAP 2007)*, Bari, Italy, December 2007.