Enlarging HPSG with lexical semantics

Toni Badia
Roser Saurí

This paper aims at enlarging the semantic treatment standardly assumed in HPSG in order to deal with several issues still not adequately solved, such as: optional verbal and nominal complements, the implication of participants and events that take part in the denotation of lexical items but are not syntactically expressable, the selection restrictions imposed by predicates to their arguments, and the non-intersective interpretation of adjectives. To this purpose we have modified and enriched the content description level of HPSG as well as its governing principle. Our semantic point of departure is the Generative Lexicon model (GL), basically because of its rich and flexible view of semantics, and its similarity to HPSG with regard to the underlying representation logic. In particular, we have taken advantage of both the GL representational aspect (that is, the multi-layered, structured conception of semantic information) and its generative dimension. The resulting proposal is implementable in LKB.

1 Introduction

Traditionally NLP systems are syntactically centered and tend to use semantics as a complement to syntactic analyses in the cases that cannot be handled by syntax alone. It is true that most theoretically oriented approaches to syntax in NLP introduce an abstract level of representation which they label as semantic. This level, however, can hardly be called semantic, if the information that is represented in it is carefully considered. There are basically two aspects that are dealt with under this heading: predicate-argument structure (which also includes modification relations) and quantification. Although quantification is an essential element in semantic analysis, we are not going to be concerned with it here, since it is not a matter of lexical semantics (but rather belongs to the structural component of semantics). Let us just mention in passing that in many cases quantification is treated only to the extent that the problems it brings about can be really avoided in parsing sentences.

Argument structure and modification, however, are both essential to syntactic analysis and central to any approach to lexical semantics. In this paper

we are interested in showing that these two perspectives can be integrated into a single approach and that the resulting system behaves better than traditional approaches. We are going to focus on HPSG because it is the theory that we know best and it uses the same underlying logic as GL, the lexical semantics system that we are now going to integrate. In addition HPSG has become one of the standards for NLP applications, so that there are now many projects that use HPSG (or HPSG-like) grammars for the syntactic processing of texts.[1] We are convinced however that nothing essential hinges on these choices: that is to say the basic ideas contained in this paper could be implemented with other syntactic theories.

In the next section we start by considering the traditional approach to argument-structure in HPSG and seeing its limitations. Section 3 introduces some of the specific data that cannot be dealt with by following the standard version of HPSG. We develop our proposed revision of the HPSG semantic treatment in section 4, and we finally apply it to the cases previously introduced (section 5).

2 Limitations of the standard approach

In HPSG argument structure and modification relations are coded, among other aspects, in the CONTENT attribute, the level that contains the restrictions to the index expressed in terms of *parametrised states of affairs* (*psoa*). The general grammatical relations that are coded at this level overcome some of the most well-known form-function mismatches.[2] Nonetheless, this level in which argument structure and modification are represented is still a direct projection of surface structure. It is certainly not a simple (or one-to-one) projection, but this is not sufficient to overcome its inherent limitations. Thus, for example, it is not easy to integrate in it argument positions that do not correspond to explicit surface positions (i.e., positions in the VALENCE lists). And this is so because the *psoa* part of content has not been designed as a full semantic representation, but simply as a deep syntactic one.

With respect to argument structure, then, the HPSG content representation is not really semantic, but simply a more abstract syntactic representation. This fact is somewhat obscured by the terminology used in this part of

[1] Some of the relevant references are: van Eynde & Schmidt [1998], Kay *et al.* [1994], and http://hpsg.stanford.edu/hpsg/lingo.html.

[2] Thus control relations are expressed by means of the coindexing of argument values in the *psoa*, so that a single element in the VALENCE list provides the content to two distinct argument positions. And passive is treated as a change in the correlation between elements in the VALENCE list and elements in the corresponding *psoa*.

the linguistic sign, which mainly derives from situation semantics. Thus attributes such as INDEX or RELATION and typed values like *parametrised state of affairs* or *individual* give the impression that what is being represented is really semantic. However, a close analysis of the phenomena studied and the treatments proposed clearly shows that the approach actually deals with the interface between syntax and semantics, not with semantics proper. Let us just pick up two examples to show this. Firstly, the richest typing of the *psoa* type that is available is the one used in binding theory [Pollard & Sag 1994:c.6]. And, secondly, the most comprehensive proposal we know of the implementation of argument structure and modification [Badia & Colominas 1998] is not aimed at representing the semantic implications of predicate classes, but only at providing a consistent typing to all complement classes.

HPSG, like other linguistic theories with a syntactic base, can adequately deal with many linguistic constructions: in general with those that rely either on strict subcategorisation relations (between the complement and its head) or on an intersective interpretation of modification. But there are constructions that can not be easily dealt with with the standard approach to subcategorisation or as intersective modifiers. In what follows we are going to examine some of these constructions: optional complements to verbs, complements to nouns in general (which are generally optional), and non-intersective modifiers.

3 Problematic data

Normally only two basic kinds of complements are distinguished: those that are strictly subcategorised by the head (sometimes called "arguments") and those that are not required for by their head –that is, modifiers. However, as has very often been recognised, this distinction is not sufficient. Firstly, it does not allow us to deal adequately with complement optionality and forces most syntactically based systems to list distinct lexical entries of verbs in order to account for their multiple realisations. Secondly, it does not allow us to represent those complements that are optional but still semantically selected by their heads, as is the case with most complements to nouns. And finally, it does not allow us to deal adequately with complements that are semantically implied even though they cannot be expressed at the surface.

Complements to verbs are often optional, but their optionality may be of different sorts. In some cases the distinction between two (or more) distinct lexical entries for the same verb might be justified, but very often this is not the case, since the presence or absence of the complement is due to syntactic and semantic properties of the sentence which have nothing to do with the

lexical semantics of the verb. This is the case, for example, of the absence of the direct object in generic sentences (1a) and object deletion structures (1b):

(1) a. La meva germana sempre dóna a ONGs
the my sister always gives to NGOs

b. Aquest noi menja molt de pressa
this boy eats very of hurry

In addition, there are cases, pointed out by Pustejovsky [1995], of complements that are clearly optional, but whose relation to the head is clearly controlled by the semantics of the verb; this is the case of the so-called *default* (*d-arg*) and *shadow arguments*(*s-arg*) respectively:

(2) a. **D-Arg**: John built the house **out of bricks**

b. **S-Arg**: Mary buttered her toast **with an expensive butter**

Complements of nouns are even more optional than compements of verbs: as a matter of fact almost every nominal complement can be omitted in some circumstance, as shown here:

(3) a. Aquesta tarda un **grup de nens** jugava a la plaça
This afternoon a group of children played in the square

b. El **grup** l' ha acceptat molt bé
The group him/her has accepted very well

(4) a. Compraré dos **fulls de cartolina**
will-buy(1st-sing) two sheets of paperboard

b. Escriu-ho en un **full**
write-it on a sheet

For complements of nouns the strategy of listing as a different lexical entry every distinct subcategorisation option is not very convincing, as there is almost no grammatical aspect that may help choose a particular lexical entry over another. This is even more problematic in languages like Catalan or Spanish in which the great majority of complements to nouns are introduced by the preposition *de*. Furthermore the choice between the objective and the subjective interpretations for complements of transitive deverbal nouns is very often not possible on simple syntactic grounds:

(5) a. l' estudi de les plantes (*the study of plants*)

b. l' avaluació dels inspectors (*the evaluation of the inspectors*)

The examples in (5) clearly show that the choice between objective and subjective interpretation strictly depends on the semantic value of those complements, since their syntactic structure is exactly the same. These facts clearly show the need to integrate full semantic information in the treatment of these complements. Further arguments can be derived from examples like those in (6), where it is shown that discourse elements can influence the interpretation of complements: as can be seen, complements that are not explicitly present in the NP or VP may serve as antecedent of an anaphoric relation or of a discourse implication. Thus the subject of *bonica* ('beautiful') (6a) and *natural* (6b) can be only the result of the acts of decorating and translating, respectively. And the use of the definite determiners *el seu* ('her') (6c) and *l'* ('the') (6d) is licensed by the "hidden" complements of *mare* ('mother') and *amanir* ('dress').

(6) a. La decoració del pont ens ha costat molt, però ha quedat tan
the decoration of-the bridge us has taken much, but has resulted so
bonica!
beautiful(fem)

b. Traduir aquest fulletó m' ha costat molt, però al final ha
to-translate this pamflet me has taken a-lot, but in-the end has
quedat molt natural
resulted very natural

c. Aquest matí ha vingut una mare. Venia a dir que el
this morning has come a mother. came(3rd-sg) to say that the
seu fill no podrà venir a l' excursió
her son not will-be-able come to the excursion

d. Hem amanit l' enciam peró l' hem llençat
have(1st-pl) dressed the salad but it have(1st-pl) thrown-away
perquè l' oli era ranci
because the oil was rancid

Modifiers can also be difficult to integrate by means of standard approaches. In some cases the adjective denotes differently depending on the context in which it appears (7).

(7) a. una biga llarga (*a long beam*)

b. una llarga tradició (*a long tradition*)

Of course, the difference here concerns the distinction between intersective and non-intersective interpretations of the adjective. But there are differences of interpretation in cases of exclusively non-intersective uses as well. Consider the adjective *ràpid* ('fast'): it usually modifies events, and

yet it can appear in expressions like those in (8) where it predicates of individuals. In section 5.2 we take into account further issues concerning this kind of adjective.

(8) a. un mecanògraf ràpid (*a fast typist*)
b. un conductor ràpid (*a fast driver*)
c. un cotxe ràpid (*a fast car*)

To deal with cases such as all of these, in the next section we modify and enrich the content description level of HPSG by integrating lexical semantics information as in GL [Pustejovsky 1995].

4 Proposed treatment

4.1 A cross-category approach to semantic information

The need for a new view of HPSG semantic level is not motivated only by the data presented here. It is actually also necessary in trying to account for two (already classic) problems in formal and computational semantics: the integration of approaches to verbal modifiers and approaches to nominals adjuncts, and the representation of the predicate structure of nominalisations. The inadequacy of HPSG in dealing with these two problems derives from its category-oriented treatment of semantics. Thus the reasonably established approach to nominal adjuncts cannot be extended to verbal modifiers since the semantic structure for verbs does not introduce any INDEX attribute to which the possible adjuncts could be linked. Similarly, nominal signs have no level where to express their predicate-argument structure, in contrast to verbs, where this level is indeed present.

We therefore argue for an homogeneous treatment across the different major syntactic categories. Such an approach is within the lines of prior work, starting in Davidson [1967]. In HPSG-related work we find Badia & Colominas [1998], where a specific treatment is proposed to deal with the above mentioned problems within HPSG, Minimal Recursion Semantics (MRS), the computational semantics framework developed by Copestake et *al.* [1998], and Sag & Wasow [1999:ch.6], which establishes a unified cross-category treatment of semantics from an HPSG-based view of grammar. Our proposal follows these works in developing a unique semantic structure for all major syntactic categories.[3] The resulting basic sign for the CONTENT level in the representation of linguistic units is as follows:

[3] As will be shown in section 4.2, nominal and verbal categories will be distinguished by the treatment of quantification.

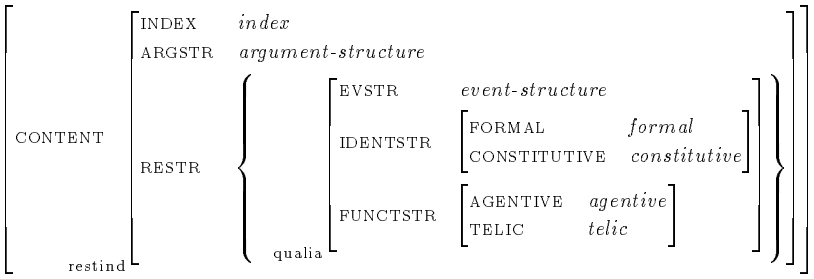

Figure 1: Basic sign for the CONTENT *level*

The modified CONTENT level integrates the INDEX and RESTRICTION (RESTR) attributes used in the description of the semantics of nominals, together with ARGUMENT-STRUCTURE (ARGSTR), which would correspond to NUCLEUS, the attribute that introduces the predicate-argument information of verbal signs in standard HPSG. We adopt here the term ARGSTR from GL since, in constrast to NUCLEUS, ARGSTR classifies the arguments in the relations expressed by the entity according to the distinction among true-, default- and shadow-arguments (cf. (2)).

A second remarkable difference between standard HPSG and our proposal concerns the appropriate type for RESTR: as can be observed, it is not a set of *psoas* but a set of *qualia* structures, each of them constituted by three different information levels: event structure (EVSTR), identity structure (IDENSTR) and functional structure (FUNCSTR). Thus, the semantic restrictions that the denoted entity has to satisfy are not constituted by a single relation or property (one per *psoa*); instead, each description level in the *qualia* structure introduces at least one relation that concerns a particular aspect of the word's meaning.

Let us now turn to the specific information in the qualia structure. For readers minimally familiarized with GL, it may appear that there is hardly any relation between the qualia structure as is stated in GL and the one introduced here. In standard GL, the qualia structure expresses lexical semantic information concerning four specific dimensions of the denoted entity: the properties that classify it within the class it belongs to (the FORMAL role), its constitutive structure (CONSTITUTIVE), its originating process (AGENTIVE) and its purpose (TELIC). Apparently, none of these dimensions directly appear as attributes of the *qualia* type in figure 1. Note, however, that two of the attributes represented, IDENSTR and FUNCSTR, group into two levels the four classic qualia roles: FORMAL and CONSTITUTIVE on the one hand, and AGENTIVE and TELIC on the other. This dual distinction is based on the lines drawn by recent work within GL, where special attention is given

to the functional qualia levels [Pustejovsky 1998, 1999]. The particular formalization that we adopt here is argued for in the following section.

A second difference between the standard qualia structure and the one shown in figure 1 is the inclusion of the EVSTR level. In GL EVSTR is an independent semantic level that represents information of the eventuality expressed by the entity. Here we include EVSTR within the *qualia* structure, which is the type appropriate for the RESTR value, because, for most kinds of predicative expressions, it conveys semantically relevant information that restricts the entity pointed at by the index and that has to be preserved as restrictive information through the processes of phrasal composition.

4.2 Differences among semantic types in the composition of phrasal units

So far we have attempted a category-unified treatment of lexical. Now we want to integrate it with the compositional dimension of meaning. We follow the treatment for quantification given in Pollard & Sag [1994]. However, some modifications have to be introduced in the subtypes of CONTENT in order to preserve the cross-category approach to basic lexical meaning. The subtypes of CONTENT in standard HPSG are the following:

$$\text{nom\_obj}\begin{bmatrix}\text{INDEX} & \textit{index}\\ \text{RESTR} & \textit{set-of-psoa}\end{bmatrix} \qquad \text{quantifier}\begin{bmatrix}\text{DET} & \textit{semdet}\\ \text{RESTIND} & \textit{npro}\end{bmatrix} \qquad \text{psoa}\begin{bmatrix}\text{QUANTS} & \textit{list-of-quants}\\ \text{NUCLEUS} & \textit{qfpsoa}\end{bmatrix}$$

Figure 2: Standard HPSG subtypes of content

The subtype *nominal_object* is the CONTENT value appropriate for non quantified nominals. It is also used as the value of RESTIND, the attribute that concerns the restrictions on the quantified index in the semantic structure of quantified expressions (*quantifier*).[4] In contrast, in the semantic structure of predicates (*psoa*), the value of NUCLEUS is a quantifier-free *psoa* (*qfpsoa*), which is a structure that expresses the relation between the predicate and its participants.

In subsection 4.1 we assumed that *restind* was the structure appropriate for the semantic representation of every major part of speech. This forces a restatement of *content* subtypes. First, the *restind* type as represented in Figure 3 is substituted for *nom_obj*. Secondly, it is also established as the value of the RESTIND attribute in the *quantifier* semantic structure. And

[4]The RESTIND value in figure 2, *npro*, is a subtype of *nominal-object*.

finally, we propose that it is also the appropriate structure to express the nuclear information in *psoa*, the semantic structure of predicates. We therefore adopt the new *restind* type as the NUCLEUS attribute (which from now on will be renamed RESTIND) in the semantic description of predicates. The three standard subtypes of *content* in figure 2 are respectively transformed as shown in figure 3:[5]

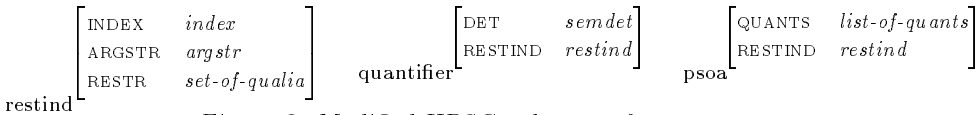

Figure 3: Modified HPSG subtypes of content

The change in the basic semantic types also requires a revision of the Semantics Principle. According to this principle, the nuclear semantic information received by the phrasal mother is taken from a different daughter depending on the *headed-phrase* type: for *head-adjunct-structures* the CONTENT value has to be token-identical to that of the semantic head daughter; otherwise, it is identical to that of the head daughter. However our proposal has introduced a substancial modification that prevents the adequate operation of the principle: the introduction in all major category of an index that expresses the kind of denoted entity (individual or eventuality). Following the standard Semantics Principle, the index value of the mother node has to be projected from the modifier and not from the syntactic head. This is now problematic in the case of head-modifier relations because, contrary to the standard HPSG version, in our version there is no correspondence between the modifier and head indices, and the one that would be (now wrongly) projected to the mother would be that of the modifier.[6] In section 5.2 we will go into this discussion in more detail. For the moment we propose to restate the Semantics Principle in order to adequately account for the composition of semantic information. Our proposal follows Sag & Wasow [1999] in assuring a 'head-driven' character

[5]For practical reasons, from now on we will use the *restind* type to represent not only the semantic structure of nonquantified nominal expressions, but also both quantified and predicative expressions, omitting the other attributes in the *quantifier* and *psoa* types.

[6]Recall that adjectives already have an index value in standard HPSG. Nevertheless, it is an index that is coindexed with the index introduced by the nominal head (i.e., an index that expresses the referent denoted by the noun being modified). The Semantics Principle, therefore, applied correctly in the composition of adjective-noun phrases since, when projecting the semantic information of the adjective to the mother node, the index value still corresponded to the nominal one.

to semantic composition in parallel with syntactic processing:[7]

In a headed phrase:

a. the RETRIEVED value is as in Pollard & Sag [1994:232]; and

b. the INDEX and ARGSTR attributes of the CONTENT value are identical to those of the head daughter, whereas the RESTR set value is composed of the union of each daughter's RESTR set.

4.3 Exploiting generative relations among word meanings

So far we have extended the semantic information level in HPSG with the representational apparatus of GL. This view of the semantic treatment of linguistic expressions is not new. An integration of GL semantic representation and HPSG syntax is proposed to some extend in Copestake [1993], Copestake & Briscoe [1996] and Johnston [1996], and more explicitly in Badia & Saurí [1998, 1999], so that a reasonably straightforward interaction between syntax and semantics become available. Furthermore, a simplified version of the standard GL representation has been used in two EU-funded projects: Acquilex and SIMPLE.[8] Common to these approaches is the fact that they do not implement the generative dimension of GL, but merely use its representational structure. However, the treatment proposed in GL relies not only on the organisation of the semantic information but also on a set of generative mechanisms which code the relations that different aspects of the meaning of words may have with one another, thus accounting for the creativity of the lexicon. Note that the generative capacity of the system is essential in order to reduce the number of lexical entries and, consequently, ambiguity in analysis. If we had a representational system that allowed us to express neatly the different aspects of word meaning but gave us no way of relating them, we would end up with a sense enumeration lexicon –more sophisticated, but still sense enumeration.

The generative devices in GL are *co-composition*, *selective-binding* and *type coercion* [Pustejovsky 1995:ch.7]. In GL implementations proposals, these devices are generally less used than the representational structure of GL, probably because they are not easy to implement. For example, the LKB used to implement the Acquilex proposals [Copestake 1993] was not powerful enough to introduce the generative mechanisms of standard GL.

[7] Since we have not modified the HPSG treatment of quantification, the part of the Semantics Principle that concerns quantification remains unaltered.

[8] The Acquilex project references are Esprit-BRA 3030 and Esprit-BRA 7315. SIMPLE is funed by EU's DG-XIII, within the LE programme.

We claim, however, that there are currently actual ways of implementing the generative capacity of the lexicon: basically, what is needed is a proper type system with multiple inheritance and enough inference capacity. If these requirements are met, either with subspecification [Markantonatou & Sadler 1998] or default inheritance [Copestake & Briscoe 1992], or with both [Lascarides & Copestake 1999], most of the devices originally contemplated in GL (and a few others) can be implemented. Implementations like the new LKB [Copestake 1998] allow for much of what is needed. In section 5 we show how our proposal, which integrates the basic HPSG sign structure with the declarative and procedural levels of GL semantics description, takes advantage of the properties of the type system in order to deal with the problems sketched in section 3.

5 Analysis of the data

5.1 Optional complements

We will first try to account for optional complements of verbal and nominal predicates, following the proposal developed in Badia & Saurí [1998]. Given their optionality, the standard HPSG treatment of obligatory complements by means of valence lists is not adequate to account for them. Nonetheless, Sanfilippo [1998] sheds some light on how they can be treated. For independent reasons he proposes that some complements can be treated as actual adjuncts from a syntactic point of view, even if they are thematically bound to the relation denoted by the head. We adopt this mechanism and represent D- and S-Args as thematically bound adjuncts, introduced as set members at the nonlocal (NONLOC) information level.

Consider first an ordinary process-result nominal such as *construcció* ('building'), a nominalisation of the verb *construir* ('build'). As stated in Pustejovsky [1995], this verb subcategorises for two obligatory complements (the agent and the result of the building process) and a third argument that expresses the material out of which the resulting entity is built. This third argument is considered a D-Arg because it is syntactically optional but participates in the logical expression of the event (cf. (2)). In the deverbal nominalisation, it is not only the 'material' argument but also the agent and result arguments that are optional. Here all three arguments are considered D-Args and are formally treated as thematically bound adjuncts à la Sanfilippo; that is, instead of being declared in the VALENCE lists, they are stated as members of the NONLOCAL set. In figure 4 we show the process reading of *construcció*. Recall that D-Args, in addition to their expression at

the NONLOCAL level, are also identified as *default arguments* at the ARGSTR.

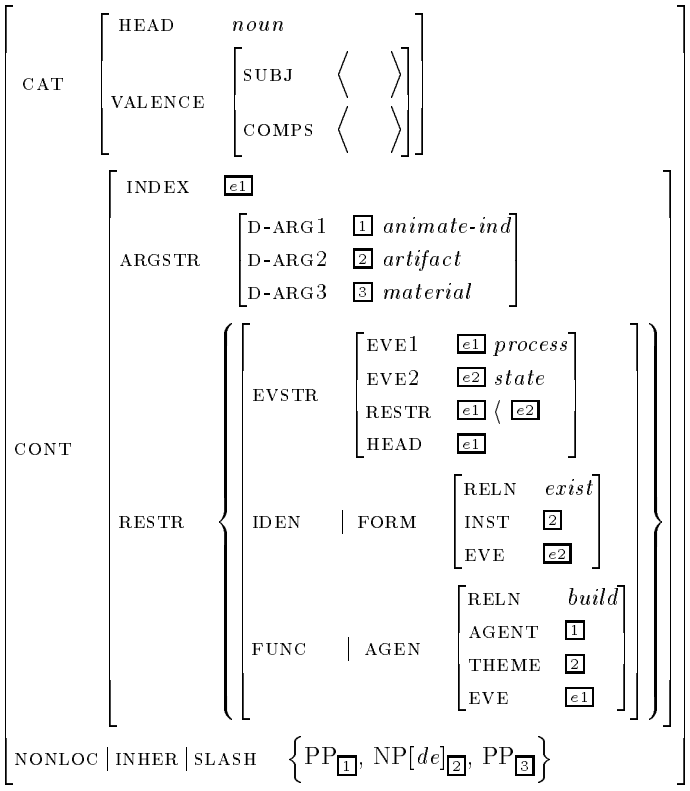

Figure 4: construcció *(process reading)*

In order to see that other types of nominals can also be treated in this way, consider for instance redescription nominalisations. They differ from standard creation predicates in that their process reading cannot express the argument denoted by the result nominalisation syntactically. Consider the noun *decoració* ('decoration'), derived from the verb *decorar* ('decorate'). As opposed to *construir*, the direct object of *decorar* does not denote an artifact that results from the decoration act, but a preexisting object being decorated. We assume that *decorar* has three arguments: one corresponding to the agent, a second one that expresses the theme (i.e., the object being decorated), and finally a D-Arg that refers to the material used in the decoration. The agent and the theme arguments are subcategorized as T-Args by the verb, but realized as thematically bound adjuncts when appearing in a nominalization.

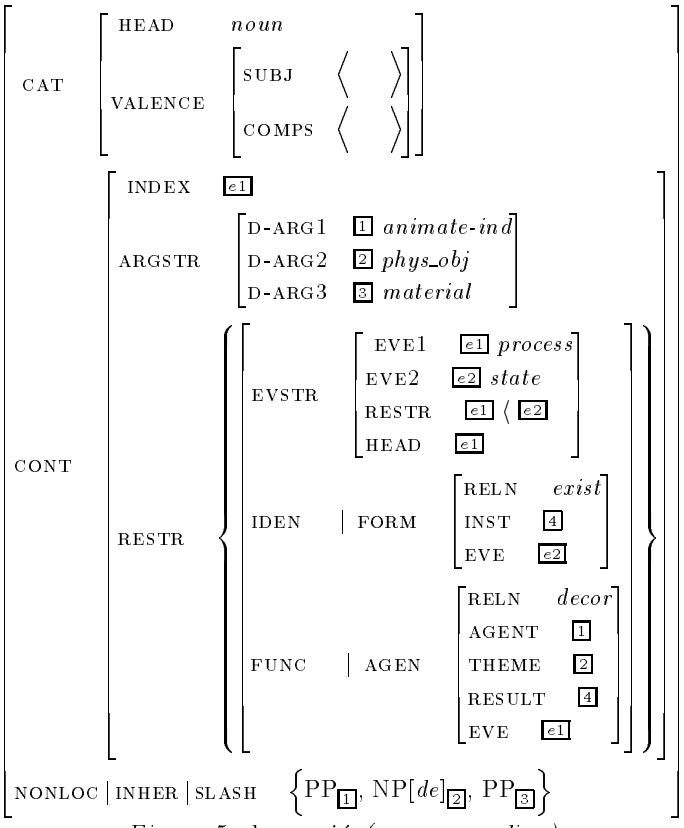

Figure 5: decoració *(process reading)*

The treatment we propose for *decoració* is shown in figure 5. Note that this adequately accounts for examples such as (6a) above, where the re-description nominal (*decoració*) expresses the process reading in the first clause, whereas in the second clause it is referred to as denoting the object resulting from the process. In the representation structure in figure 5 the reference to the result is allowed by the argument of the formal qualia role, which is not bound by any argument in the ARGSTR just because it can never be syntactically realised as an argument. Thus the argument structure acts as interface between the rich semantic representation (i.e., the set of qualia structures in the value of RESTR) and the surface mechanisms that license predicate complements (VALENCE lists and NONLOCAL set): only semantic arguments that may be syntactically realised are present in the argument structure (either as T-Args or as D- or S-Args). In addition, the rich semantic structure of the sign allows us to express semantically implied arguments, and thus provides a treatment for semantically motivated discourse factors like the ones shown in (6) above.

The proposal outlined here provides an appropriate treatment for creation and redescription predicates in general. Instead of the ad hoc solution adopted in Badia [1997] for a particular case (*traducció*, 'translation'), our treatment derives from a general and systematic approach to the semantic structure of predicates and their nominalisations. This approach explains why anaphoric reference to the unexpressed result argument of the process reading is possible. This treatment is also applicable to the verbs *menjar* and *amanir* in examples (1b) and (6d) above. In addition, the rich semantic information we use allows us to deal in a natural way with the distinction between subjective and objective complements shown in (5). Other types of nouns with semantically implied (optional) complements can be similarly treated; for example, non-deverbal nouns that express a relation with another entity, like nouns denoting sets or partition (*grup* 'group' in (3) and *full* 'sheet' in (4)) or relational nouns (*mare* 'mother' in (6c)).

5.2 Non-intersective modification

We now turn to non-intersective (nominal) modifiers, which, as pointed out above, also demand a revision of the standard HPSG semantic treatment. The problems illustrated by the adjective *ràpid* ('fast') in (8) above are two. On the one hand, the adjective presents a non-intersective interpretation: it is generally an eventuality predicate but here it modifies individual-denoting nouns. On the other hand, it denotes differently ('who types fast', 'who drives fast', 'that can be driven fast') depending on the noun it combines with [Bartsch 1985], although there is indeed a semantic core that is common to all three instances of *ràpid* –that is, the property of being fast of a given event.

Larson [1998] explains similar non-intersective cases by adapting Davidson's event analysis, originally developed for adverbs, into the semantic structure of the nominal expressions. His proposal, particularly focussed on agentive nouns like *dancer* or *typist*, provides good insight into the problem but leaves some aspects unresolved, such as the pervasivity of event modification in nominals. Interestingly, however, Pustejovsky's GL approach offers an adequate and systematic treatment of these facts. If we assume that *ràpid* ('fast') is an event predicate, then we can argue it triggers an event interpretation for the noun it modifies. This can be done by applying the selective-binding mechanism, which forces the adjective to predicate over the qualia level which is adequate to its selectional restrictions (i.e., an event), instead of predicating over the whole entity. Thus, when modifying *mecanògraf* ('typist'), *ràpid* predicates of the process of typing, the event

encoded at the telic level of the semantic structure of the noun, whereas with *conductor* ('driver'), *ràpid* predicates of the 'driving' event.

GL, therefore, provides an elegant treatment of the non-intersective use of adjectives that predicate of events. It is also general enough to explain their apparent sense variation depending on the noun they appear with. These advantages are mainly due to two factors. One, the distinction between individual- and eventuality-denoting adjectives; and two, the introduction of a structured multi-layered semantic level for the description of the content of nominals (and other categories). Nevertheless, these two factors entail the twofold revision of the standard HPSG semantic treatment of modifier-head relations discussed in section 4: the assumption that there are two distinct sorts of indices forces a revision of the Semantics Principle, whereas the complexity in the semantic content level of nouns makes clear the need for a highlighting mechanism of semantic information. The next two subsections willbriefly discuss these revision, while the third subsection develops the issue of the implementation of GL's generative mechanisms in a real typed system.

5.2.1 Index information and its processing through composition processes

The issue we will focus on here was introduced in section 4.2, where a modification of the Semantics Principle was put forward. Here we want to revisit it in connection to the examples in (8) and in light of the recently proposed distinction between individual- and eventuality-denoting indices. Once the distinction is assumed, it turns out that certain adjective-noun phrases cannot be adequately handled by standard HPSG semantics treatments.

Phrases composed of an individual-modifying adjective and a nominal head are no problem for HPSG. Already in Pollard & Sag [1994], the index (of type individual) introduced by the adjective is coindexed with the index of the nominal head and, by the standard Semantics Principle, it is projected to the semantic structure of the mother node. Similarly, the case of event-modyfing adjectives that combine with typically event-denoting nouns (such as *construcció* ('building') or *decoració* ('decoration') is not a problem either as long as eventuality types of indices are available: the eventuality index introduced by the adjective is coindexed with the index the nominal head and is projected to the mother node by means of the Semantics Principle much in the same way as individual indices. Problems appear, however, in cases of non-intersective modification such as those in (8). Here the (eventuality) index of the adjective is of a different sort from the (individual) index of the

nominal head. Besides, the index projected to the mother node would be the modifier's eventuality index, altough the entity being referred to by the whole phrase (for instance, *a fast typist*) corresponds to an individual. To overcome this inadequacy (within a head-driven approach to semantic composition), we have put forward the modification of the Semantics Principle in section 4.2, which simply forces the head to project its semantic information to the mother node.

However, this analysis interferes with another important aspect of classic HPSG semantics: the specific nature of the INDEX feature. Indeed, in the case of nouns indices introduce the kind of entity referred to by the word. This is, however, not so with adjectives, where the index value corresponds to the entity (individual or eventuality) being modified, instead of expressing the property-denoting character of the adjective. Therefore our proposed modification also brings about a change in the semantic structure of adjectives and other modifiers, the index of which will introduce the type of entity denoted by the adjective (i.e., a state –a kind of eventuality). This revision has positive repercussions for the general treatment of adjectives: they now express a part of nuclear information to which their possible complements and modifiers can be linked, and some information about the head they modify.

5.2.2 Semantic prominence among qualia roles

Let us now turn to the issue of the extension of the semantics level with a multi-layered structure derived from the qualia and eventive structures in GL. As already seen, this representation level (combined with the selective-binding mechanism) allows for an adequate treatment of non-intersective modifiers such as those in (8). However, there is still one unsolved issue: there are at least two event values (at the agentive and telic level) in the qualia structure of all nouns, and it is not clear how event-selecting adjectives manage to choose between them. The adjective *ràpid* ('fast') provides examples of this:

(9) a. un mecanògraf ràpid (*a fast typist*)

b. un cotxe ràpid (*a fast car*)

(10) a. un pastís ràpid (*a quick cake*)

b. una construcció ràpida (*a fast building*)

Ràpid selects the telic quale of the noun when modifying *mecanògraf* ('typist') or *cotxe* ('car'), (9); that is, it selects the information about the

goal process related to the denoted entity. But when it combines with *pastís* ('cake') or *construcció* ('building'), *ràpid* selects the agentive level (10), which conveys the information concerning the genesis process of the entity.

It is our intuition that not all nominal lexical entries have their qualia structured in the same way. That is to say, every nominal class has a particular quale role which is more prominent than the others. For instance, instrumental and agentive nouns (such as *knife* and *typist*, respectively) are characterised by the prominence in their telic quale; whereas in result nominalisations (such as *building*) and nouns like *statue* the most prominent event level is the agentive quale. For the sake of brevity we will not discuss this issue any further here (see [Badia & Saurí 1999] for detail). What mainly interests us here is how to implement the enrichment of the information in the qualia structure of nominals in order to indicate the particular quale role that is prominent in each case (where applicable). As an example, figure 6 shows the entry for *ganivet* ('knife'), a noun that highlights the telic role:[9]

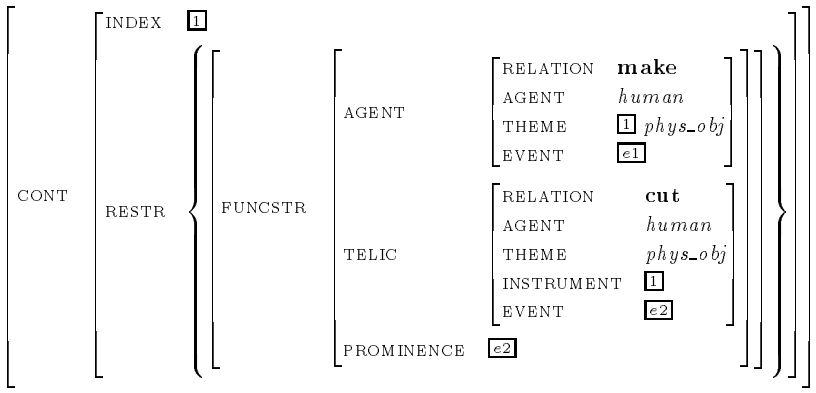

Figure 6: ganivet *('knife')*

Prominence highlights a particular piece of the noun's semantic information. The four GL qualia roles are subject to tensions and oppositions between them, which are manifested through analogous pieces of information: between the formal and the constitutive qualia, on the one hand, as the roles that express identity properties of entities, and between the agentive and telic information, on the other hand, as eventuality-referring levels. Although for reasons of space this account is very roughly sketched here, this fact is what grounds the division into two pairs of the four classic GL qualia roles: one named *identity structure*, which concerns the identity of the entity referred to by the noun, and a second one which concerns its functionality

[9]For reasons of space, from now on we only show the relevant semantic levels.

(*functional structure*). Within this picture, prominence is, then, the feature that expresses the strongest role within each of these two basic relations.

The notion of prominence turns out to be necessary in adjective-noun composition processes; particularly in those cases where the adjective (individual- or eventuality-modifying) is underspecified as to the particular qualia role that it selects for. In these cases the adjective predicates of the prominent quale in the noun.[10] Take figure 7 as an example; the requirements imposed by a simple eventuality-modifying adjective like *ràpid* ('fast') on non-predicative nouns (such as *cotxe* 'car', and *pastís* 'cake') would be expressed basically as shown:

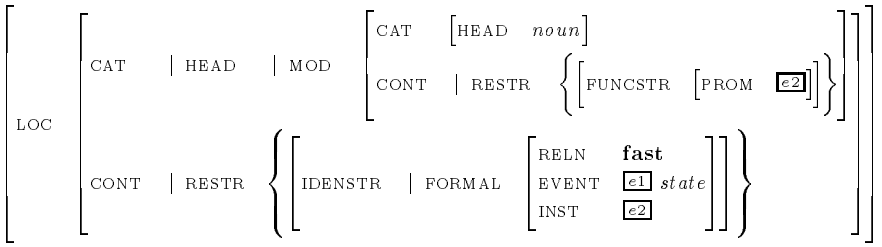

Figure 7: ràpid *('fast')*

The adjective's content level states that *ràpid* predicates the property of being fast of an eventuality which in turn corresponds to the prominent eventuality in the FUNCSTR of the modified noun. This structure represents the information that *ràpid* should contain in order to allow for a non-intersective use (as in examples (8)). But *ràpid* is actually a modifier of eventuality-denoting nouns (such as *construcció* ('building') and *decoració* ('decoration') in their process reading) and therefore selects for nouns with an index value of type eventuality. It is precisely in the cases where the nominal head does not comply with this requirement that a non-intersective interpretation of the adjective is allowed by means of the selective-binding mechanism. The issue arising at this point is how to implement within a real typed feature system the selective-binding mechanism used to explain the non-intersective use of certain adjectives.

5.2.3 Implementing GL mechanisms

The three generative mechanisms in Pustejovsky [1995] are adopted as operative devices to explain certain recursive relations observed between words.

[10] The fact that nouns present two different prominent qualia (one in IDENSTR and the other in FUNCSTR) is not a problem: the former is established among individual-type indices, whereas the latter is chosen from eventualities.

However a closer look reveals that they actually constitute the description of processes that occur between word meanings. Taking as an example the case of *ràpid* previously analised, we see that the selective-binding mechanism is an abstraction of the non-intersective interpretation process. Therefore there is a need to establish a formal treatment of this process and others that are subject to comparable relations. For our implementation we use the (new) LKB system [Copestake 1998], a grammar and lexicon development environment which is specifically designed for the use of typed feature structures with underspecification and multiple default inheritance. Such a flexible and robust platform allows us to implement the GL generative mechanisms by simply exploiting the expressiveness of the type system, instead of having to view them as extra processes that apply to the lexicon.

Given that non-intersective uses of adjectives are secondary to the common intersective ones, we establish a partial hierarchy for eventuality-modifying adjectives. It consists in a first general type (*eventuality_modifier_adjective*), which represents the intersective use of adjectives, and a second type (*subeventuality_modifier_adjective*), which is a subtype of the former type and represents the fact that these adjectives may predicate over a "subeventuality" within the nominal head. Since part of the information of both types is incompatible (basically, the semantic level in the noun structure to which the eventuality index belongs) we need some overwriting mechanism. To implement it we take benefit of YADU ('Yet Another Default Unification'), the default representation proposal in Lascarides & Copestake [1999] which is effectively integrated into LKB. In YADU, types are represented by means of bipartite structures (typed default feature structures (TDFSs)) of the form *Indefeasibe/Tail*: *Indefeasible* is a simple typed feature structure that expresses what is indefeasible, whereas TAIL, which specifies the defeasible information, consists of a set of pairs where the first member of the pair is an atomic feature structure (a single path or equivalence) and the second one is a type.

The partial YADU hierarchy needed to account for both intersective and non-intersective uses of eventuality-modifying adjectives is as shown in figure 8.[11] The indefeasible information stated in the general supertype is completely subsumed by the subtype. The difference between both TDFSs is in the Tail, where the supertype asks for an eventuality-denoting noun,

[11] Due to space limitations, we will not represent the first members of each pair in the tail set as an atomic feature structure. Instead, we integrate all of them in a unique, non-atomic feature structure –this is why there is just one pair in both tails. In addition, we have abbreviated some of the (already abbreviated) attribute names: C | H | M stands for CAT | HEAD | MOD, whereas R stands for RESTR.

whereas its subtype selects for the prominent eventuality in the functional structure of the noun –the other content level where eventuality-type indexes are stated. Note that the information in the subtype TDFS basically corresponds to that stated in figure 7.

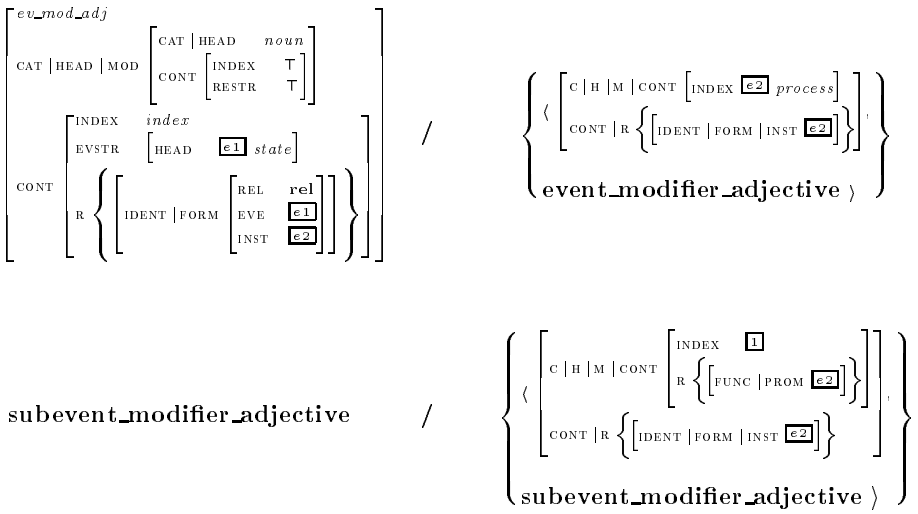

Figure 8: Partial type hierarchy for ràpid*-like adjectives*

6 Conclusions

In this paper we have shown that a strictly syntactically oriented approach is clearly inadequate to deal with (both verbal and nominal) complement optionality and non-intersective modification. A basic requirement for a more adequate treatment is to have available a semantic information level that, although connected, is independent of the syntactic level. Furthermore, this content level should be based on a rich and robust conception of semantics. This allows us to deal with several problematic issues in a natural way: the implication of participants and events that take part in the denotation of lexical items but are not syntactically expressable, the selection restrictions imposed by predicates on their arguments, and non-intersective uses of adjectives.

We have introduced several modifications into the standard HPSG content level by adopting a cross-category approach to semantics. Our grounding semantic model has been GL, from which we have taken adopted with its representational dimension and its generative capacity. We have also devel-

oped an implementation of one of the GL generative mechanisms in a proper type system (the new LKB by Copestake [1998]), with multiple inheritance and default unification.

References

[Badia, T., 1997] 'Predicative structure of nominals in HPSG'. In *Syntax and Semantics of Predicates* (1997 Conference of the TLS). Austin. University of Texas.

[Badia, T. & C. Colominas, 1998] 'Predicate-Argument Structure'. In Van Eynde & Schmidt (eds.) *Linguistic specifications for typed feature structure formalisms.* Luxembourg, EC.

[Badia, T. & R. Saurí, 1998] 'The Representation of Syntactically Unexpressed Complements to Nouns'. *COLING-ACL'98. Workshop on the Computational Treatment of Nominals*, p: 1-10. Montréal.

[Badia, T. & R. Saurí, 1999] 'Semantic disambiguation of adjectives in local context: a generative approach'. In P. Bouillon & E. Viegas *Description des Adjectifs pour les Traitements Informatiques.* Workshop. TALN'99. Corsica.

[Bartsch, R., 1985] 'The Structure of Words meanings: Polysemy, Metaphor, Metonimy'. In F. Landman & F. Veltman, Eds., *Varieties of Formal Semantics.* Dordrecht: Foris.

[Copestake, A., 1993] *The Compleat LKB.* Acquilex-II deliverable, 3.1. ms. CCL.

[Copestake, A., 1998] *The (New) LKB system.* http://hypatia.stanford.edu/ aac/lkb.html (updated: Dec 10,1998).

[Copestake, A. & T. Briscoe, 1992] 'Lexical operations in a unification-based framework'. Pustejovsky, J. & S. Bergler (eds.)*Lexical Semantics and Knowledge Representation. Proceedings of the ACL SIGLEX Workshop on Lexical Semantics and Knowledge Representation.* Berkeley. Berlin, Springler-Verlag: 109-119.

[Copestake, A. & T. Briscoe, 1996] 'Semi-productive polysemy and sense extension'. In Pustejovsky, J. & B. Boguraev (eds.) *Lexical Semantics. The problem of polysemy.* Oxford, Oxford University Press.

[Copestake, A., D. Flickinger & I. Sag, 1998] *Minimal Recursion Semantics. An Introduction.* CSLI, Stanford University. ftp://csli-ftp.stanford.edu/linguistics/sag/mrs.ps.gz

[Davidson, D., 1967] 'The Logical Form of Action Sentences'. In N. Rescher, Ed., *The Logic of Decision and Action.* Pittsburgh. University of Pittsburgh Press: 81-120.

[Johnston, M., 1996] 'Semantic underspecification in lexical types: capturing polysemy without lexical rules'. *Acquilex Workshop on Lexical Rules, 1995.* Cambridge.

[Kay, M., Gawron & Norvig (eds.), 1994] *Verbmobil, a translation system for face-to-face dialog* Standford, CSLI.

[Larson, R.K., 1998] 'Events and Modification in Nominals'. *Proceedings of the Semantics and Linguistics Theory.*

[Lascarides, A. & A. Copestake, 1999] 'Default Representation in Constraint-based Frameworks'. *Computational Linguistics*, vol. 25, 1: 55-105.

[Markantonatou, S. & L. Sadler, 1998] 'Lexical generalisations'. Van Eynde & Schmidt (eds.) *Linguistic specifications for typed feature structure formalisms.* Luxembourg, EC.

[Pollard, C. & I. Sag, 1994] *Head-driven Phrase Structure Grammar.* Standford CA: CSLI.

[Pustejovsky, J., 1995] *The Generative Lexicon.* Cambridge MA: The MIT Press.

[Pustejovsky, J., 1998] *The Semantics of Lexical Underspecification.* ms. Dept. of Computer Science, Brandeis U.

[Pustejovsky, J., 1999] *Type Construction and the Logic of Concepts.* ms. Dept. of Computer Science, Brandeis U.

[Sag, I. & T. Wasow, 1999] *Syntactic Theory: A formal Introduction.* Stanford, CSLI.

[Sanfilippo, A., 1998] 'Thematically bound adjuncts'. In Balari & Dini (eds.) *Romance in HPSG.* CSLI. Stanford.

[van Eynde & P. Schmidt (eds.), 1998] *Linguistic specifications for typed feature structure formalisms.* Luxembourg EC.

Toni Badia is professor at Universitat Pompeu Fabra and head of the research unit in Computational Linguistics of the Institut Universitari de Lingüística Aplicada (IULA) of the same University. Universitat Pompeu Fabra. La Rambla, 30-32. 08002 Barcelona. Catalonia, Spain. toni.badia@trad.upf.es.

Roser Saurí is a researcher of the research unit in Computational Linguistics of the Institut Universitari de Lingüística Aplicada (IULA). Universitat Pompeu Fabra. La Rambla, 30-32. 08002 Barcelona. Catalonia, Spain. roser.sauri@trad.upf.es