Tilburg University

Default inheritance in an object-oriented representation of linguistic categories

Daelemans, W.M.P.; De Smedt, K.; de Graaf, J.

Publication date:
1991

Document Version
Publisher's PDF, also known as Version of record

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):
Daelemans, W. M. P., De Smedt, K., & de Graaf, J. (1991). *Default inheritance in an object-oriented representation of linguistic categories*. (ITK Research Report). Institute for Language Technology and Artificial Intelligence, Tilburg University.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

CBM

CBM
R

8409
1991

31

UNIVERSITY
LIBRARY
UNIVERSITEIT
BRABANT

ITK

RESEARCH
REPORT

ITK Research Report
November 15, 1991

**Default Inheritance in an
Object-Oriented Representation
of Linguistic Categories.**

W. Daelemans, K. De Smedt and
J. de Graaf
No. 31

ISSN 0924-7807

©1991. Institute for Language Technology and Artificial Intelligence,
Tilburg University, P.O.Box 90153, 5000 LE Tilburg, The Netherlands
Phone: +3113 663113, Fax: +3113 663110.

DEFAULT INHERITANCE IN AN OBJECT-ORIENTED REPRESENTATION OF LINGUISTIC CATEGORIES

Walter Daelemans

ITK

University of Brabant

P.O.Box 90153, 5000 LE Tilburg,

The Netherlands

walter@kub.nl

Koenraad De Smedt and Josje de Graaf

NICI

University of Nijmegen

P.O.Box 9104, 6500 HE Nijmegen,

The Netherlands

desmedt@nici.kun.nl

ABSTRACT

We describe an object-oriented approach to the representation of linguistic knowledge. Rather than devising a dedicated grammar formalism, we explore the use of powerful but domain-independent object-oriented languages. We use default inheritance to organize regular and exceptional behavior of linguistic categories. Examples from our work in the areas of morphology, syntax and the lexicon are provided. Special attention is given to multiple inheritance, which is used for the composition of new categories out of existing ones, and to structured inheritance, which is used to predict, among other things, to which rule domain a word form belongs.

1 INTRODUCTION

During the last few decades, research in knowledge representation and research in computational linguistics have been getting closer to each other, but in two different ways. On the one hand, the frame-based and object-oriented knowledge representation languages used in AI have widened their grasp on linguistic knowledge: not only domain knowledge has been 'framed', but also syntax, morphology, phonology, and the lexicon (Daelemans, 1987, 1988; De Smedt, 1984, 1990). On the other hand, dedicated linguistic formalisms have been enriched by ideas coming from established work in knowledge representation. The incorporation of inheritance in unification-based formalisms (e.g. Shieber, 1986) is an example of such an enrichment.

We argue that it is better to refine and further tailor a general (but sophisticated) computer language to the needs of the linguistic domain, than to design a new language, reinventing the object-oriented wheel, from scratch. We also wish to embed linguistic competence in a more comprehensive theory of cognition as much as possible, thereby allowing maximal theoretic generalization. Proponents of Word Grammar (Hudson, 1984), among others, have set the same goal. In this context, we argue that the application of AI languages to linguistic knowledge deserves more attention. Specifically, the mechanisms for inheritance, encapsulation, and polymorphism, which are common in frame-based and object-oriented programming are essentially domain-independent. They are eminently suited to the representation of many kinds of knowledge, including linguistic knowledge. In this article, we will concentrate on several aspects of default inheritance to illustrate this point, and we will use selected parts of our work in morphology, syntax and the lexicon of Dutch as an example domain.

1.1 Language as an open system

What we do in defining a knowledge representation language is to make an abstraction over a class of representational structures and introduce a syntactic mechanism to express that abstraction. The resulting new primitives will manage the complexity of knowledge so that programs will be more understandable, modularity will improve, etc. Thus there is a practical gain. But the relevance of postulating new representational primitives goes beyond mere productivity concerns. They state generalizations about the representational structures used by processes in intelligent systems (Steels, 1978). As with all empirical generalizations, it may not be possible to absolutely prove that it is adequate or valid, but it may be possible to find cases demonstrating why a generalization is notationally adequate, for example from a viewpoint of parsimony.

Default inheritance in taxonomies of object classes or types is such a representational primitive. It refers to the ability to generalize and specialize mental concepts, and to reason by virtue of prototypes while keeping the system open to exceptions. This form of reasoning

pervades much of our common sense knowledge but is also a high-level mechanism for the symbolic manipulation of important concepts in a scientific domain. Fikes and Kehler (1985) point out that representations based on hierarchically ordered, structured objects “capture the way experts typically think about much of their knowledge, provide a concise structural representation of useful relations, and support a concise definition-by-specialization technique that is easy for most domain experts to use.” This is also true for the way linguists generally think.

The software engineering advantages in using object-oriented programming (modularity, conciseness, reusability) are well-documented and apply *a fortiori* to the design and implementation of natural language processing systems, where hierarchical reasoning with defaults is necessary for a practical and realistic representation of linguistic knowledge. Language is an open system. The development of extensible and adaptable natural language processing systems crucially depends on a knowledge representation paradigm within which generalizations are effectively exploited (Jacobs, 1985).

This is not to say that inheritance is only a software engineering tool. Especially in a lexicalized grammar, the avoidance of redundancy is mandatory, while the possibility to incorporate exceptions must be left open (Flickinger, 1987). We are trying to model linguistic knowledge the way linguists typically work. Many grammars, especially traditional ones, are implicitly organized by means of abstraction over linguistic categories. The hierarchical relation between grammar concepts may emerge in the organization of a grammar book: one finds a chapter on *The noun*, which in turn has a more specialized section on *The proper noun*, etc. While we think this is basically the right approach, it is of course necessary to make the nature of the inheritance relations in the hierarchy much more explicit. This can be achieved by developing formal theories of default inheritance and proposing algorithms for their implementation (see Daelemans et al., 1992, for an overview of the use of inheritance in natural language processing). In short, the achievement of a hierarchical theory of language must be supported by a knowledge representation framework which incorporates primitives for hierarchical reasoning.

1.2 Inheritance as a multi-purpose mechanism

An important feature of object-oriented languages is that they provide some kind of mechanism for objects to inherit their structure and behavior from other ones. The application of inheritance can be seen from different points of view:

- 1 *Specialization*. From a conceptual point of view, inheritance allows objects to be specializations of more general objects. In this way a *specialization hierarchy* is produced, which corresponds closely to the hierarchy of ‘*is a*’ relations in a semantic network. For example, *vowel* and *consonant* might be specializations of *phonological segment*. A hierarchy of grammar concepts might contain objects for *word* and its specializations *noun*, *verb*, *preposition*, etc.

2 *Combination.* Another kind of inheritance, *multiple* inheritance, represents the behavior of an object as a combination of the behaviors of two or more other objects. This is often done if an object needs to integrate knowledge from different sources or perspectives. For example, *John* may be a male patient; thus the behavior of *man* and that of *patient* is combined. Composition will often consist of the addition of a few special features (for example *patient*) to a more general category (in this case *man*); the secondary object whose features are ‘added’ is sometimes called a *mixin*. A *transitive compound strong verb* could be defined as a combination of the object *verb* with the mixins *transitive*, *compound*, and *strong*.

3 *Stepwise refinement.* A program can be constructed by first modeling the most general concepts in the application domain, and then dealing with special cases through more specialized objects. The programmer is not so much concerned here with the construction of a taxonomy but uses refinement as a programming methodology. Stepwise refinement by specialization can be compared with the well-known methodology of stepwise refinement by decomposition (Wirth, 1971). It is significant that the effort of defining an object is *proportional* to the extent in which it differs from other objects. Thus, refinement is not only useful as a programming methodology, but it can also be thought of as a general cognitive mechanism, for it reflects a *principle of least effort*.

4 *Avoiding redundancy.* From the point of view of data storage, inheritance is a way of knowledge sharing. A piece of information which is necessary in many objects needs to be stored in only one object, where others can access it. Avoiding redundancy in this way reduces memory requirements. It also improves modularity, because the shared knowledge only needs to be updated in one place. Again, the principle of system economy is thought of as a general cognitive principle and not just a software engineering strategy.

5 *Predictions about new objects.* When new objects are created in the course of the computation, inheritance relations can be used to predict their default behavior. For example, if we say that the mother of *person* inherits from *woman*, we predict that if a certain object is the mother of a specific person, it will—at least by default—be a woman. Such prediction can be specified in *structured inheritance* relations. Structured inheritance offers many opportunities for the representation of linguistic knowledge and will be dealt with in more detail below.

It is important to mention here that our view of inheritance is always based on *defaults* rather than absolute and irretractable statements. If we say “birds fly”, then we mean “birds typically fly” and we have no problems to accept and handle exceptions—for instance, ostriches and penguins—adequately. Similarly, if we say “nouns are countable”, then there may still be a special class of uncountable nouns. All inherited knowledge only holds in so far as it is not overruled by knowledge in the inheriting object. Consequently, our view of inheritance incorporates an *implicit blocking* of defaults in view of more specific knowledge.

We conduct our work in the frame-inspired object-oriented languages CommonORBIT (De Smedt, 1987, 1989) and KRS (Van Marcke, 1987) which are based on the notion of prototypes. For examples we will use the syntax of CommonORBIT. Whereas in some other

object-oriented languages, every object must belong to some class (or type, or flavor), this is not the case in CommonORBIT. An object in CommonORBIT can be a prototype (or model, or proxy) for any other object. This view of inheritance uses 'is-like' or 'shares-with' relations, which are more general than the meaning associated with 'is-a' relations in languages which involve typing.

2 A HIERARCHICAL MODEL OF DUTCH MORPHOLOGY

We have developed and implemented a model of Dutch morphology which generates structured, phonologically and orthographically specified word forms from their bases in the lexicon. The model, which covers Dutch nouns, verbs and adjectives, consists of the following modules: (1) a hierarchical lexicon with objects representing simple unstructured words (base forms), (2) a morphological component consisting of a hierarchy of morphological categories and associated rules, and (3) a phonological component consisting of a hierarchy of phonological categories and rules. The three modules are not independent in their representation; as will be shown below, they are only different locations in the topology of a single large lexical hierarchy of linguistic concepts. The system is embedded in a larger natural language generation¹ system, which provides base forms with features such as singular, third person, etc. On the output side of the morphological component, a set of phonological rules (for instance, assimilation) operate on the morphological representation to yield a fully specified phonetic string, and spelling rules yield an orthographic string.

2.1 The lexicon as a hierarchy

The morphological knowledge forms part of a larger object-oriented lexicon architecture in which maximal use is made of the information combination, default reasoning, and sharing possibilities of multiple default inheritance. Lexical base forms have associated with them isiosyncratic information (their lexical representation), and inherit as leaf nodes from different objects in different subhierarchies (semantic, morphological, syntactic, orthographical). In this way, both procedural knowledge (i.e. rules) and declarative knowledge (e.g. features) become available to the base by default inheritance.

In Figure 1, part of such a hierarchy is shown. The base form (*to*) *pay* is defined as inheriting from *transitive*, *regular verb*, and *commercial action*. The hierarchy represents surface semantic, syntactic, and morphological knowledge. Through default inheritance, the base form gains access to the syntactic knowledge that it is a main verb, that its object is an

¹ One may wonder how suitable these methods are for applications other than generation. We will not address this question in depth here, but we admit that it is not straightforward to rearrange the system so that it is suitable for parsing. Although generation and parsing are assumed to operate largely on the same kinds of structures, it is not possible to simply run the programs described here 'in reverse'.

NP with accusative case, and that its subject is an NP with nominative case. The semantics inherited includes the knowledge that the agent of the action referred to by ‘to pay’ is the referent of the subject, and the patient is the referent of the object. By virtue of the inheritance link to *regular verb*², a number of word formation rules (methods) become available to the base form, in which new word forms are created that are defined in their proper place in the hierarchy. That way, lexical information becomes available through inheritance to these derived forms as well.

Figure 1

Integration of knowledge in a hierarchical lexicon

The uniform representation of lexical knowledge enables the definition of the interaction between different sources of knowledge at a high level of abstraction (e.g. the definition of semantic agents as syntactic subjects), while at the same time allowing for exceptions and sub-regularities (e.g. subjects of passives are patients). In the remainder of this article, we will mainly focus on the organization of the morphological subhierarchy within this lexical architecture.

² Morphologically speaking, *to pay* is seen as a regular verb despite superficial spelling changes in its past tense and past participle.

2.2 Word formation as a recursive process

Dutch polymorphemic word forms may be derived from simpler forms by means of prefixing, for instance, *her+doe* (redo), by suffixing, for instance, *groen+ig* (greenish), or by a combination of both prefixing and suffixing, for instance, *ge+werk+t* (worked). This process covers both derivation and inflection. In addition, words may be compounded, for instance, *lucht+haven* (airport) or can undergo simultaneous compounding and affixing, for instance, *drie+kopp+ig*³ (three-headed). As can be expected, polymorphemic word forms may themselves be the base for other word formation processes as if they were simple morphemes, which makes word formation a recursive process. It is easy to assign a structure to polymorphemic words which reflects the ordering of the various recursive derivations and compoundings by which they are generated. For example, the Dutch word *ge+her+structur+eer+d+e* (restructured) can be represented as a bracketed structure where each successive recursive layer of the derivation is represented as a concatenation of a prefix, a base, and a suffix, where the prefix or the suffix may possibly be empty (" "):

(" " (ge+ (her+ (" " structur +eer) " ") +d) +e)

This representation, which we call the *lexical representation*, is computed by means of the list concatenation of the lexical representations of a (possibly empty) prefix, a base, and a (possibly empty) suffix. The CommonORBIT code of the prototypical wordform contains a procedure for creating such a list:

```
(DEFOBJECT WORD
  (PREFIX (A 0-MORPHEME)) ;default = empty
  (BASE :IF-NEEDED #'IDENTITY) ; default = the word itself
  (SUFFIX (A 0-MORPHEME)) ;default = empty
  (LEXICAL-REPRESENTATION :IF-NEEDED
    #'(LAMBDA (SELF)
      (LIST
        (LEXICAL-REPRESENTATION (PREFIX SELF))
        (LEXICAL-REPRESENTATION (BASE SELF))
        (LEXICAL-REPRESENTATION (SUFFIX SELF))))))

(DEFOBJECT 0-MORPHEME
  MORPHEME
  (LEXICAL-REPRESENTATION " "))
```

The first DEFOBJECT form defines an object *word* with aspects *prefix*, *base*, *suffix* and *lexical-representation*. The second DEFOBJECT form defines the empty morpheme. Aspects of type :IF-NEEDED contain functions that compute a value only when needed. For more details on the syntax of CommonORBIT, we refer to De Smedt (1987). The lexical

³ Our morphology abstracts from spelling adjustments, e.g. reduplication of the final consonant in *kopp*. These adjustments are carried out by a separate spelling module which is active on the same level as the phonological component.

representation of a morpheme is a string consisting of phonological segments⁴ and boundaries, supplemented by diacritics for stress, etc. The lexical representation of word is computed as a list structure containing other lexical representations, i.e. those of its prefix, base and suffix. A phonological string can be derived from such a list by removing the brackets and concatenating the contained strings. It should be realized, however, that the lexical representation, as we use it, abstracts from regular phonological variation. The ultimate phonetic string is determined only after general phonological rules (for instance, assimilation) have operated on it.

2.3 Objects as a uniform representation

Our work uniformly represents all linguistic categories, such as phonemes, morphemes, etc. as structured objects. Several basic tenets of object-oriented programming are relevant to linguistic representation. One of these is the principle of encapsulation, the localization of structure and behavior in the object to which it conceptually belongs. An object can be seen as a collection of properties and (potential) behaviors. Furthermore, the identity of an object is important⁵. This identity can be used, for instance, to relate stems of complex word forms such as *inherit* and *heritage* (see also Russell, Carroll and Warwick, 1990, for a treatment of separable verbs along these lines).

This implies that in our approach to morphology, not strings are the basic units (as in two-level morphology or generative phonology), but structured morphological objects. The phonetic representation, which takes the form of a string, and the lexical representation, which takes the form of a list, are only two of many attributes associated with word form instances. Similarly, string concatenation is only one aspect of morphological processing, which involves object creation and multiple inheritance of properties. This approach makes it easier and more natural to handle discontinuous morphemes, morphological processes that have no effect on word form (e.g. type conversion), and multi-word lexical entries (idioms).

With respect to hierarchical representation, our work follows De Smedt (1984), who accounts for some regularities and exceptions in the inflectional morphology of Dutch verbs by employing default inheritance in a hierarchy⁶ as depicted in Figure 2.

⁴ For ease of reading, we will often use spelling rather than pronunciation in this article.

⁵ For a discussion of the importance of object identity in the representation of linguistic knowledge, we refer to Van der Linden et al. (1988).

⁶ This example is also discussed in Gazdar (1987).

Figure 2

A partial hierarchy of Dutch verbs (from De Smedt, 1984)

The top node of this hierarchy, *weak verb*, is the most regular kind of verb and therefore is the prototype object for all other verbs. Actually, this object inherits from an even more general object, *word*, and so on. The inflectional behavior of the weak verb is represented in a number of aspects which each compute a stem, prefix, and suffix for each of its inflections. All weak verbs have the object *weak verb* as a prototype. For example, *werken* (to work) is an object which inherits its inflectional behavior directly from *weak verb*. There is a specialization representing a subclass, *mixed verb*, which is partly regular but has an exceptional past participle suffix. This contradicts the specific information associated with its prototype for this particular form, which is thus overruled. Again, there are a number of verbs which inherit their morphological behavior from *mixed verb*, there is a subclass, and so on. This representation provides a non-redundant and generalization-capturing account of Dutch verb inflections.

It is significant that the hierarchies in Figures 1 and 2 contain lexical items as well as morphological categories. Thus, there is no strict separation between the lexicon and the body of morphological knowledge. Instead, there is a smooth transition from the most general categories to the most specific ones—individual words. This is a direct consequence of the uniform representation of linguistic knowledge as objects. A uniform framework has the advantage that it allows for the same general principles to be applied to a variety of knowledge units (Jacobs, 1985).

3 MULTIPLE INHERITANCE AS COMPOSITION

Specialization is but one facet of the inheritance mechanism. The other is the definition of new objects as a composition of several prototypes. This presupposes the handling of *multiple* inheritance relations. This possibility can be exploited in at least three ways.

3.1 Combined objects as mixins

One way to exploit composition is the creation of categories that combine information from different knowledge sources. This will often amount to the addition of a few special features to a more general category. For example, a *nominative plural NP* can be seen as an object that inherits from *nominative*, *plural*, and *NP*. The prototype *NP* is the most substantial category, whereas *nominative* and *plural* are secondary objects whose features are ‘added’; they are sometimes called *mixins*. We will avoid the many thorny issues in multiple inheritance (see for example Touretzky, Horty & Thomason, 1987), but nevertheless we must raise the question of how conflicts between contradicting knowledge in the composing objects can be resolved. It is clear that mixins are often meant to have priority over the defaults in the more general categories. In the case of a *nominative plural NP*, the mixins *nominative* and *plural* have priority over the prototype *NP*, which may be accusative and singular by default. We will call the relative ordering of composing objects *local* or *definitional precedence*.

However, we must make sure that definitional precedence does not violate hierarchical precedence. For example, suppose that *rondrijden* (to ride about, to tour), a Dutch compound verb, inherits from both *compound* and *verb*. The prototypes *compound* and *verb* are defined separately, so that the knowledge encapsulated in them can be reused for different combinations, e.g. *compound noun*, *compound adjective*, etc. Suppose, furthermore, that we specify that the knowledge in *compound* has definitional precedence over that in *verb*. Any knowledge in *compound* will precede that in *verb*. In principle, we could implement this as a depth-first search in the hierarchy (left-to-right in Figure 3).

Figure 3

Multiple inheritance with a common prototype

When combining knowledge from different sources, it cannot be excluded that the different sources share a common prototype, for instance, the object *word* in Figure 3. A depth-first search in the hierarchy will reflect the priority relation between *compound* and *verb*, but will not do justice to the specialization relations. The common prototype must not be considered before more specialized objects in the hierarchy. This can be formulated as the following principle, which has also been advocated by Ducournau and Habib (1987):

Specialization vs. Multiplicity:

Inheritance must follow the specialization partial order; therefore, in any case the specialization relation excels the local (definitional) precedence of prototypes.

Following this principle, *compound* and *verb* will have precedence over *word* in the example of Figure 3.

3.2 Multiple inheritance as biased semi-compositionality

Seen from a different perspective, the combined operation of the principles of definitional precedence and specialization yields the kind of 'biased' semi-compositionality that is common in descriptions of natural language phenomena. One prototype usually plays the role of *head* while the others play the role of *modifiers* with different priority. Furthermore, the information which is compositionally inherited from the prototypes can be overruled by locally specified information to express subgeneralizations and exceptions.

The use of composition as consisting of a head and modifiers can be exemplified by the semi-compositional nature of compounds such as *blackboard*. This compound could be defined as inheriting from *board* and *black* (in that definitional precedence order), biasing the composition of information toward inheritance of syntactic and morphological features from *board*. However, the combination is not only biased, but also only semi-compositional, because in this case, semantic information is not always inherited from the semantic representations of the parts (blackboards may be green by default).

3.3 New categories as mules

A third use of composition creates new categories whose behavior is 'in between' that of two or more others. Think of such a new category as a mule which is the young of a donkey and a horse. We will give a more elaborate example in the domain of the morphology of verbs, in order to show this use of multiple inheritance. In Dutch, there are in fact more kinds of strong and half-strong verbs than those dealt with in section 2.3. Some half-strong verbs are different in that they have a vowel change in the past participle; this kind will be called *half-strong verb 2*. A third kind is partially weak and partially strong, but in exactly the opposite way: they have a strong past tense (with vowel change), but a weak past participle; let us call this kind *half-strong verb 3*. The new hierarchy is depicted in Figure 4.

Figure 4

Objects for Dutch half-strong verbs

The strong verbs in Dutch can also be divided into two kinds, one which exhibits a vowel change in the past participle and one which does not. This offers the opportunity to use multiple inheritance for object composition, because the behavior of the strong verbs can be found distributed among the various half-strong verbs, as shown in Figure 5. This representation is especially powerful because the definitions for the strong verbs consist *only* of the combination of the objects they inherit from, without any additionally specified knowledge. The multiple inheritance principle (specialization vs. multiplicity) makes sure that more specialized prototypes have priority over more general ones. Thus, for example, *strong verb 1* inherits from both *half-strong verb 1* and from *half-strong verb 3* before the more general knowledge in *weak verb* is considered. Thus, the defaults in *weak verb* are effectively blocked. In this hierarchy, there is no conflict between branches because the main classes of half-strong verbs are opposite and thus complementary; they each provide differing specific information which does not contradict each other. However, the principle of specialization vs. multiplicity is crucial.

Figure 5

Revised hierarchy of Dutch verbs

4 STRUCTURED INHERITANCE

We now want to devote special attention to the application of structured inheritance in morphology, because this inheritance mechanism, though very powerful, has received little attention in the computational linguistics literature.

4.1 Structured inheritance as an automatic mechanism

Structured inheritance is a mechanism in frame-based and object-oriented representation which models a slot after one higher in the hierarchy. When an object (or frame) inherits from another one, the fillers of its aspects (or slots) will automatically inherit from corresponding fillers in the higher level object. Structured inheritance is central in the frame-based language KL-ONE (Brachman & Schmolze, 1985) and in CommonORBIT.

Let us first consider a simple example in a non-linguistic domain. Suppose that we want to express the knowledge that the mother of a person is normally a woman who is (at least by default) not a virgin. We could represent this by creating an object for the concept *person* and defining a mother aspect which is filled by an object representing the prototypical mother, i.e. a 'non-virgin woman'. The inheritance relations are graphically represented by means of arrows in Figure 6. Notice that there is a stacking of defaults, one in *woman* and one in the

mother of *person*. The corresponding code in the object-oriented language CommonORBIT is as follows (where each expression is followed by the result of its evaluation):

```
(DEFOBJECT PERSON
  "The mother of a person is, by default, a woman who is,
  by default, not a virgin."
  (MOTHER :OBJECT (A WOMAN
                    (VIRGIN? NIL))))
#<object PERSON>

(DEFOBJECT WOMAN
  "A woman is a person of the female sex."
  PERSON
  (SEX :VALUE 'FEMALE))
#<object WOMAN>
```

Figure 6

Structured inheritance: a non-linguistic example

Structured inheritance is a way of using this representation to infer some property of a specific mother. When we ask for the mother of a specific person, say Olivia, we obtain an object which inherits from the prototypical mother of *person*. Figure 7 graphically illustrates the following CommonORBIT code:

```
(DEFOBJECT OLIVIA WOMAN)
#<object OLIVIA>

(MOTHER 'OLIVIA)
#<the MOTHER of OLIVIA>

(SEX (MOTHER 'OLIVIA))
FEMALE

(VIRGIN? (MOTHER 'OLIVIA))
NIL
```

Figure 7

Structured inheritance at work

The identity of the newly made object is important. For each object which inherits from person, a new mother object is created, such that every person object is provided with its own unique mother.

Summing up, structured inheritance has been described as the ability to “... preserve a complex set of relations between description parts as one moves down the specialization hierarchy” (Brachman & Schmolze, 1985:177). In systems providing structured inheritance, the inheritance mechanism is not limited to simply sharing or copying a value, but it models an object and the network of all its associated objects after one higher up in the hierarchy, thereby potentially eliminating considerable redundancy. Since structured inheritance is transitive, knowledge from various levels can be combined. Finally, since defaults are handled as in ordinary inheritance, it remains perfectly possible to represent exceptions. For example, it would be straightforward to represent an object with a virgin mother in CommonORBIT as follows:

```

(DEFOBJECT JESUS PERSON
  ((MOTHER VIRGIN?) T)
  #<object JESUS>

```

4.2 Derivation rules as generic functions

The basis of the morphological system is formed by objects representing morphosyntactic categories. As is usually done, the category label is represented as a feature. For example, an *adjective* is a *word* with a *category* feature which has value *adj*. This prototype for *adjective* is defined in CommonORBIT as follows:

```

(DEFOBJECT ADJECTIVE
  WORD
  (CATEGORY 'ADJ))
#<object ADJECTIVE>

```


Word formation is characterized as the computation of a derived form (inflection, derivation)⁷ from a base form. Each base form belongs to one or more classes, each corresponding to part of the domain of a derivation rule. The rule itself is represented as a *generic function* whose behavior depends on the class of the objects in its domain. Hence, the conditional part of a rule can be *distributed* over a number of categories.

Morphological rules are attached to the objects in their domain, i.e. they are defined as procedural aspects (or methods) of objects representing morphological categories. By way of example, the rule generating a derivation with the suffix *+ig* ('ish'), for example *groen+ig* ('greenish') is attached to a category representing the domain of this rule:

```
(DEFOBJECT ADJ-WITH-IG
  ADJECTIVE
  (IG (AN ADJECTIVE
    (BASE :IF-NEEDED (SELF)
      (BASE (WHERE SELF = IG))
      (SUFFIX (AN IG-SUFFIX))))))
#<object ADJ-WITH-IG>

(DEFOBJECT IG-SUFFIX
  MORPHEME
  (LEXICAL-REPRESENTATION "+IG"))
#<object IG-SUFFIX>
```

This DEFOBJECT form is read as follows. *Adj-with-ig* is a category designating the prototypical adjective which has a (possible) derived form with *+ig*, stored as the value of an aspect *ig*. Using a type theory terminology, we could read *ig* as a function applying to objects of type *adj-with-ig*, with a return value of type *adjective*. The derived form, i.e. the return value of the *ig* function, is an adjective with a *base* and a *suffix*. The *base* of the derived form is the same as the *base* of the object of which the derived form is derived (*where self = ig*). The *suffix* of the derived form is always an *ig-suffix*, i.e. a morpheme with the lexical representation *+IG*.

Other, more specific objects may inherit from the prototype *adj-with-ig* to access its method to perform the derivation. Each individual adjective inheriting from the prototype will have its own derived form. The set of objects inheriting from the prototype can thus be viewed as the domain of the rule.

Again, following De Smedt (1984), these morphological classes are placed in a hierarchy, where new classes are formed by means of specialization and combination. The leaves of this hierarchy are the base words, i.e. the lexicon. For example, the Dutch adjective *groen* ('green') is in the domain of the *ig*-derivation. It is represented as an object which inherits from the category *adj-with-ig*; in addition it has a phonologically specified base:

⁷ From the purely morphological point of view, we do not make a principled distinction between derivation and inflection, calling both processes by the name of derivation instead.


```
(DEFOBJECT GROEN
  ADJ-WITH-IG
  (BASE (A MORPHEME
    (LEXICAL-REPRESENTATION "#Grun"))))
#<object GROEN>
```

This object will be dynamically modeled upon *adj-with-ig* by means of structured inheritance, and so the filler of the *ig* aspect will inherit from that in the prototype *adj-with-ig*, so that *groen* will have its own derived form with *+ig*. Since it was specified in the prototype that the base of the derived form is same as the base of the form of which it was derived, we have sufficient knowledge to establish that the derived form for this particular adjective is *groen+ig* (or *#Grun+ig*, phonologically). This adjective is created only when the generic function *ig* is applied to the object *groen*, as shown below.

```
(LEXICAL-REPRESENTATION (IG 'GROEN))
(#Grun +IG)
```

4.3 The range of one rule as the domain of another rule

Structured inheritance is especially useful when we want to specify that an object in the range of a rule is in the domain of another rule. For example, from *groen+ig* a comparative *groen+ig+er* ('greenisher') can be derived. Other possible orders of recursive word formation may be ungrammatical and must be ruled out (e.g. **groen+er+ig*). Using the functional metaphor, it makes sense to represent the result of the *ig*-derivation as an object which is in the domain of the comparative rule. This is schematically represented in Figure 8.

Figure 8

Words in the domains of rules

This knowledge can be represented in an object-oriented way as the constellation of objects in Figure 9, where gray dots represent objects automatically created by means of structured inheritance.

Figure 9

Structured inheritance creates derived forms and constrains them

The definition of *adj-with-ig* is thus adapted so that it constrains the result of the *ig*-derivation so that it is an adjective which may form a comparative:

```
(DEFOBJECT ADJ-WITH-IG
  ADJECTIVE
  (IG (AN ADJECTIVE-WITH-COMP ;may form comparative
    (BASE :IF-NEEDED (SELF)
      (BASE (WHERE SELF = IG))
      (SUFFIX (AN IG-SUFFIX))))))
#<object ADJ-WITH-IG>
```

The form *groen+ig+er* is then created by a double derivation, i.e. an application of the generic function *comparative* to the result of the *ig*-derivation:

```
(IG 'GROEN)
#<an ADJECTIVE-WITH-COMP>

(LEXICAL-REPRESENTATION (COMPARATIVE (IG 'GROEN))
  ((#Grun +IG) +@r))
```

4.4 Valency reduction as a hierarchical phenomenon

The objects in the range of a derivation tend to have a smaller morphological valency than those in its domain. This phenomenon is known as *valency reduction*. For example, *groen* can have both a derivation with *+ig* (*groen+ig*) and a comparative with *+er* (*groen+er*), while the *ig*-form itself cannot undergo another derivation with *+ig* (**groen+ig+ig*). Nor is it possible to reverse the order of the suffixes, i.e. **groen+er+ig*. Thus the base form *groen* has the highest valency, whereas its derived forms have ever diminishing valency. Valency reduction is easily represented by means of inheritance relations as depicted in Figure 10.

Figure 10

Valency reduction and the inheritance hierarchy

The highest object in the hierarchy has the lowest valency. Higher valency objects are formed by specializations which successively place the category in yet another domain. The lowest objects in the hierarchy are the base forms which have the highest valency. At the same time, the organization of the hierarchy is one of the factors which account for the order of morphemes in a derived form.

Summing up, we have presented a minimally redundant way to organize a lexicon in a hierarchical way such that we account for the ordering constraints on suffixes. We have presented a classification of words in terms of the domains of the rules to which they belong. This rule-oriented organization is not the same kind of hierarchy as the morpheme-oriented organization which was presented in sections 2.3 and 3.3. A lexical object can conceivably be linked to both hierarchies.

4.5 Competing productive derivations and exceptions

Dutch has several possible suffixes for the formation of plural nouns. In addition to some non-productive paradigms, there are two competing productive patterns, one with the suffix *+n/en* and one with *+s*. Both have an open domain, where the domain of *+s* is marked by a number of conditions, and *+n/en* applies otherwise. The domains are therefore not marked by separate classes, but by a condition, here summarized as *conditions-for-s*:

```
(DEFOBJECT N-WITH-PLURAL
  N
  (PLURAL (A N
    (BASE :IF-NEEDED (SELF)
      (BASE (WHERE SELF = PLURAL)) ;same as singular
    (SUFFIX :IF-NEEDED (SELF)
      (COND ((CONDITIONS-FOR-S SELF)
        (AN S-MORPHEME))
        (T (AN EN-MORPHEME)))))))
```

Since structured inheritance is a form of default inheritance, exceptions specified in the lexicon will override the inherited information. For example, the Dutch noun *zee* ('sea')

would normally get a plural with +s, but this must be overruled, because its plural is *zeeën*. This exception can be specified in a straightforward way, as shown in the CommonORBIT code and depicted in Figure 11.

```
(DEFOBJECT ZEE
  N-WITH-PLURAL
  ((PLURAL SUFFIX) (AN EN-MORPHEME)))
#<object ZEE>

(DEFOBJECT EN-MORPHEME
  MORPHEME
  (LEXICAL-REPRESENTATION "+@n"))
#<object EN-MORPHEME>
```

Figure 11

A default suffix predicted by structured inheritance is overruled

Similarly, the base of a plural noun form is generally equal to that of the singular, but there are some exceptions like *stad/sted+en* ('city/cities', with a kind of umlauting similar to German *Stadt/Städte*). We define the object *stad* so that the base of the plural is a morpheme with lexical representation "*sted*". The exceptional nature of this lexical entry is depicted in Figure 12, where it can be seen that the default base of the plural noun is overridden.

```
(DEFOBJECT STAD
  N-WITH-PLURAL
  ((PLURAL BASE) (A MORPHEME
    (LEXICAL-REPRESENTATION "sted"))))
#<object STAD>
```

Figure 12

A default base predicted by structured inheritance is overruled

In a final and more intricate example, infixes are added by overriding a simple base with a structured one. Whereas the default suffix for female counterparts of person names is *+e*, the female counterparts of verbal derivations ending on *+er* are derived by infixation, for example *werk+st+er* ('female worker') from *werk+er* ('worker'). By giving the former the structure *((werk+st)+er)*, we reinterpret the infixation as a suffixing operation on the base. To this end, we define both the base and the suffix of the derived form as exceptions to the defaults, as shown in Figure 13.

Figure 13

Infixing as the overruling of a default base with a structured base

In the schematic representation of Figure 13, the symbol '<-' denotes a pointer to the object to the left of the current one, i.e. the object in which the current object is an aspect filler. Using such paths of pointers, it can be seen that the *base* of the *fem* of the *er* of *werk* has the same *base* as *werk* itself.

5 UNIFICATION AND INHERITANCE

Of the list of useful functions of inheritance which was presented in 1.2, *constraint checking* seems to be lacking. This is of course one of the main functions of *unification* in unification-based formalisms (UBFs). However, it is possible to obtain the same functionality by letting multiple inheritance notify failure when contradicting information is inherited from different sources (multiple monotonic inheritance). This approach cannot be combined with default reasoning, however (unless when using two different inheritance mechanisms). Another approach to unification consist of the definition of a *unify* method that destructively merges objects to be unified.

5.1 Feature structures as objects

It is not difficult to see the analogy between UBFs and object-oriented formalisms. Feature structures can be represented as structured objects, and features as slots (aspects). Complex values are again objects (making the structures recursive), and atomic values are objects without slots, as well as other atoms such as symbols, strings and numbers. Paths are a series of pointers from one object to another. The following shows a CommonORBIT object definition and an equivalent feature structure:

```
(A FEATURE-STRUCTURE
  (CATEGORY 'NP)
  (PLURAL '-')
  (NOMINATIVE '+))
[category = ,NP,plural = ,-,nominative = ,+ ]
```

Reentrancy (or feature sharing) is usually represented in UBFs by means of special labels; in object-oriented formalisms it is simply token identity. Using templates in UBF boils down to using inheritance. One way of using templates is the definition of mixins to hold specific feature values, e.g., an object *singular* which contains only the value '-' for the feature *plural*:

```
(DEFOBJECT SINGULAR
  (PLURAL '-'))
```

Using such mixins as abbreviations, the CommonORBIT definition of a feature structure can be much more concise. The following definition:

```
(A SINGULAR NOMINATIVE NP)
```

again describes the feature structure given above. Figure 14 shows part of the hierarchy which contributes to this object.

Figure 14

Example hierarchy of objects as feature structures

The necessity of being able to work with types or with default inheritance (if only to reduce redundancy) in UBFs has become increasingly clear and has given rise to a number of monotonic or non-monotonic extensions or complementations, e.g. *overwriting* in PATR-II, Kaplan's *priority union*, Shieber's *add conservatively*, *datatyping* in UCG (Moens et al., 1989), CLE *sortal restrictions* (Alshawhi et al. 1989), *default unification* (Bouma, 1990), feature structure and slot-filler *typing* in HPSG (Pollard and Sag, 1987), DATR as a default reasoning formalism complementary with a UBF (Evans and Gazdar, 1989), etc. Default inheritance is incorporated in most object-oriented languages, so we take the reverse approach and represent features structures as objects, which allows them to use all the default reasoning machinery of the object-oriented formalism.

5.2 Segment Grammar

Segment Grammar (SG) is a unification-based formalism which is especially suited to incremental syntactic processing. Originally proposed by Kempen (1987), it has been further worked out and implemented by De Smedt & Kempen (1991) for the incremental generation system IPF. The basic units of SG are *syntactic segments* which represent individual syntactic relations between two categories. The top node of a segment in the context of a syntactic tree is called the *root* and the bottom node the *foot*. Typical examples are NP-head-N, representing the relation between an NP and its head noun, and S-subject-NP, representing the subject relation between an S and an NP.

Syntactic segments join together by means of a variant of the unification operation described above to form larger structures. The recursive⁸ unification operation succeeds if the values of all these features in both objects match in one of the following ways: (1) If the values are both objects, then the objects are unified; the unification (if successful) becomes the new feature value; (2) otherwise, the disjunctive values are interpreted as sets and their intersection is computed; the intersection (if not empty) becomes the new feature value. If unification succeeds, then the two objects are merged into one, and the new feature values are stored into this one object, as well as all other information which was present in both objects.

SG has been implemented in CommonORBIT by uniformly representing all important grammar units—syntactic segments, syntactic categories (phrases, words) and syntactic features—as structured objects. Inheritance allows the grammar to be extended easily by creating segments as specializations or combinations of other ones. The specialization hierarchy of segments exploits multiple inheritance. E.g., knowledge common to both segments S-subject-S and S-subject-NP is stored in a general segment S-subject-*. Knowledge common to all subordinate clauses is stored in *-*-S. By way of example, part of the segment hierarchy of a Dutch grammar is shown in Figure 15.

Figure 15

Part of the hierarchy of syntactic segments

The structured inheritance mechanism in CommonORBIT establishes inheritance relations between the root of a segment and that in its prototype, and likewise between the foot of a segment and that of its prototype. For example, the lexical segment for *water* inherits from NP-head-N. The structured inheritance relations which are automatically established are depicted in Figure 16. From this hierarchy, it can be seen that the noun *water* inherits from the N at the root of the NP-head-N segment and can thus, for example, obtain knowledge about the typical surface positions of a head noun (not explicitly shown in the example), which again eliminates considerable redundancy.

⁸ The full recursive power of unification in CommonORBIT is not strictly necessary in Segment Grammar, because grammatical relations, which require the recursion, are not represented in the same way as grammatical features, which are not recursive.

Figure 16

Structured inheritance in syntactic segments

The question of how this knowledge is used in a sentence generation task, falls outside the scope of this paper. For more details, the reader is referred to De Smedt and Kempen (1991).

6 MULTI-ATTRIBUTES

In unification-based approaches, knowledge is retrieved from feature structures by a process akin to function application. Applying the name of an attribute to a feature structure returns the value for that attribute in that feature structure⁹. In the previous section, we pointed out the similarity of this to the object-oriented approach (instance variables are encapsulated within a single object). However, multiple default inheritance can be generalized in a way that is impossible to achieve in a unification-based approach. We introduce *multi-attributes*, inspired by multimethods in CLOS (Keene, 1989), as a means to associate attributes with *combinations* of objects. This provides us with an extremely powerful, minimally redundant, and notationally adequate way to describe linguistic generalizations in cases where many prototypes interact to determine a linguistic decision.

Consider as an example German weak adjectives. The choice of the suffix is determined by case, number and gender. The data in Table 1 are taken from Zwicky (1985). There are two possible suffixes: *+en* and *+e*. The suffix *+en* is the default. Direct (accusative or nominative) singular weak adjectives get *+e*, and an exception to this is constituted by the accusative masculine singular form, which gets *+en*. Figure 17 shows the prototype

⁹ We must add a note here about the functional nature of feature structures. Feature structures are usually seen as functions which map features onto values; in fact, they are sometimes called *functional* structures for this reason. In a sense, an object-oriented representation in CommonORBIT does the reverse in the sense that features are *generic* functions which map feature structures onto values. Some object-oriented languages (e.g. FLAVORS: Weinreb & Moon, 1980) do actually implement objects as functions.

hierarchies playing a role in the problem. We have a situation here where the choice of the suffix cannot be predicted from any one of the morphological prototypes representing adjective classes, syntactic features, or suffixes, but only from the cooccurrence of a number of them.

Table 1
German weak adjective endings

	SING			PLUR		
	MASC	NEUT	FEM	MASC	NEUT	FEM
NOM	+e	+e	+e	+en	+en	+en
ACC	+en	+e	+e	+en	+en	+en
GEN	+en	+en	+en	+en	+en	+en
DAT	+en	+en	+en	+en	+en	+en

Figure 17
Hierarchies for adjectives and syntactic features

The only possible way to describe this with the type of inheritance discussed so far, would be to create ad hoc prototypes *direct-singular* as a subtype of *weak-adjective*, and *accusative-masculine* as a subtype of *direct-singular*. As multi-attributes can be associated with aggregates of prototypes, flexibility increases considerably. The following code (using the syntax of CLOS multi-methods) shows how the assignments would have to be formulated.

```

(DEFMETHOD SUFFIX ((ADJ WEAK) CASE NUMBER GENDER)
  '+EN)

(DEFMETHOD SUFFIX ((ADJ WEAK) (CASE DIRECT) (NUMBER SING) GENDER)
  '+E)

```



```
(DEFMETHOD SUFFIX ((ADJ WEAK) (CASE ACC) NUMBER (GENDER MASC))
  '+EN)
```

```
(SUFFIX BREIT NOM SING MASC) > +E
```

With three assignments, we have described all the relevant data (involving 24 possible combinations of feature values for each weak adjective). The power of multi-attributes derives from the flexibility they provide in accessing arbitrary regions in a multi-dimensional space formed by different feature hierarchies, while at the same time allowing default reasoning. They also allow the use of several independent hierarchies where otherwise one deep and tangled hierarchy would have to be used.

Another example of the expressive power of multi-attributes is the description of German separable verbs. In Russell et al. (1990), a default inheritance treatment of this phenomenon is presented. We will adopt most of the prototypes and general organization used there, and restrict ourselves to demonstrating how multi-attributes can be used to improve notational adequacy.

German separable verbs are a subtype of prefixed compound verbs in which the prefix is a bound morpheme both when the verb is untensed and when it is the head of a verb-final clause. This situation is marked by a feature *INV = no*. In the solution of Russell et al., this is described by associating three *variant sets* of feature equations with the class separable corresponding with the situations where the verb is untensed, tensed and *INV = no*, and tensed and *INV = yes*. In the latter case, the prefix of the verb is empty (i.e. a null morpheme). In Figure 18, a hierarchy similar to the one in Russell et al. is shown. The main difference is that two small independent hierarchies for tensedness and inversion are introduced.

Figure 18

Hierarchies for German separable verbs

The influence of syntactic context on the separateness of the prefix from the verb can now be easily described using the following rules, again using the syntax of CLOS.

```
(DEFMETHOD PREFIX (VERB INV FIN)
  (A 0-MORPHEME) )
```

```

(DEFMETHOD PREFIX ((VERB PREFIXED-VERB) INV FIN)
  (PREFIXFORM SELF))

(DEFMETHOD PREFIX ((VERB SEPARABLE-VERB) (INV INVERSION)
  (FIN INFINITIVE))
  (A 0-MORPHEME))

```

The first multi-attribute definition states that all verbs whatever their syntactic context lack a prefix. The second default states that prefixed verbs have as prefix the prefix-form of their prefix, again whatever the syntactic context. Finally, the third multi-attribute states that separable prefixed verbs when they are tensed and at the same time occur as head of a non-verb-final clause have no prefix (i.e. the prefix is a null morpheme).

7 CONCLUDING REMARKS

We have shown that natural language is a knowledge domain which benefits from the use of inheritance. We have demonstrated how inheritance can be applied for multiple purposes. We have treated the use of inheritance not only for specialization, but also for combination of defaults by means of multiple inheritance, and we have shown how structured inheritance can predict the behavior of new objects which are created by other objects. Our examples have shown how traditional linguistic notions, such as exceptions and blocking, can be transparently modeled by means of hierarchical reasoning with defaults. All of this can be programmed in an object-oriented language. The elimination of redundancy, which is achieved by using inheritance, should not merely be seen as a way to save memory resources, but primarily as the basis for abstraction in linguistic theory.

Several similar approaches to the representation of linguistic knowledge as structured objects can be found in the literature. In this paper, we have concentrated on some specific kinds of inheritance—multiple and structured inheritance—and have shown how they can be relevant for linguistics. For a general overview of the use of inheritance in natural language processing, we refer the reader to Daelemans, De Smedt and Gazdar (1992). Inheritance can be put to work in two different ways. Either a grammar is expressed in a domain-independent knowledge representation which provides inheritance as a basic mechanism, or inheritance is added to a special grammar formalism. We argue that the former approach is more efficient, because it avoids duplication of the inheritance mechanism for different domains. It is far easier to create a UBF in CommonORBIT than to simulate CommonORBIT in a UBF, simply because CommonORBIT already provides all necessary primitives to represent feature structures. Furthermore, it is almost implicit in the use of a domain-independent representation that linguistic and non-linguistic knowledge share a common basis for their representation. Indeed, this uniformity of representation makes it possible to bring out and exploit commonalities in different cognitive domains. We see no reason why the operation of default inheritance in a linguistic domain would be different from that in other cognitive domains.

One problem with an object-oriented approach (shared by all symbolic approaches to natural language processing), is the fact that there are many alternative ways a domain can be modeled. The designer has to make explicit choices about which object types, attributes, and taxonomies to choose. Furthermore, once designed, models are fairly inflexible and rigid. The main concern for future research in inheritance-based natural language processing (as well as all other inheritance-based reasoning) should therefore be the development of techniques for automatically acquiring and adapting hierarchies of prototypes.

A second area for future research is the extension of object-oriented formalisms with notions from other network-based paradigms. For example, object-oriented formalisms could be extended with activation levels for prototypes. This activation level can then be used to dynamically determine their precedence during multiple inheritance. This is useful to model contextual influences on word sense disambiguation, and misclassification (e.g. overgeneralization) in morphology. Attributes can be given an activation level as well, which makes a dynamic definition of object equality possible. Activation of a particular object or attribute may be the result of contextual bias, be relative to frequency, or to any other notion of salience. Incorporating the notion of activation into an object-oriented model supports hybrid symbolic-associative models, combining the strengths of both approaches.

REFERENCES

- Bobrow, R.J. & Webber, B.L. 1980. Knowledge representation for syntactic/semantic processing. In: *Proceedings of the First Annual National Conference on Artificial Intelligence* (pp. 316-323). Stanford University.
- Bouma, G. 1990. Defaults in Unification Grammar. *Proceedings of the ACL*, Pittsburgh.
- Brachman, R.J. & Schmolze, J.G. 1985. An overview of the KL-ONE knowledge representation system. *Cognitive Science* 9, 171-216.
- Brewka, G. 1989. Nonmonotonic logics—a brief overview. *AI Communications* 2, 88-97.
- Daelemans, W. 1987. *Studies in language technology: an object-oriented model of morphophonological aspects of Dutch*. Ph.D. dissertation, University of Leuven, Department of Linguistics.
- Daelemans, W. 1988. A model of Dutch morphophonology and its applications. *AI Communications* 1 (2), 18-25.
- Daelemans, W., De Smedt, K. and Gazdar, G. 1992. Inheritance in Natural Language Processing. *Computational Linguistics* (forthcoming).
- De Smedt, K. 1984. Using object-oriented knowledge-representation techniques in morphology and syntax programming. In: O'Shea, T. (ed.) *Proceedings of the 6th European Conference on Artificial Intelligence* (pp. 181-184). Amsterdam: Elsevier.
- De Smedt, K. 1987. Object-oriented programming in Flavors and CommonORBIT. In: Hawley, R. (ed.) *Artificial Intelligence programming environments* (pp. 157-176). Chichester: Ellis Horwood.

- De Smedt, K. 1989. *Object-oriented knowledge representation in CommonORBIT*. Internal Report 89-NICI-01, University of Nijmegen, Nijmegen Institute for Cognition research and Information technology (NICI).
- De Smedt, K. & Kempen, G. 1991. Segment Grammar: a formalism for incremental sentence generation. In: Paris, C.L., Swartout, W.R. & Mann, W.C. (eds.) *Natural language generation in Artificial Intelligence and Computational Linguistics* (pp. 329-349). Boston / Dordrecht / London: Kluwer Academic Publishers.
- Ducournau, R. & Habib, M. 1987. On some algorithms for multiple inheritance in object-oriented programming. In: *Proceedings of ECOOP'78 (Bigre+Globule 54)*, 291-300. Paris: AFCET.
- Evans, R. and G. Gazdar. 1989. The semantics of DATR. *Proceedings of the EACL*, Manchester.
- Fikes, R. & Kehler, T. 1985. The role of frame-based representation in reasoning. *CACM* 28, 904-920.
- Flickinger, D. 1987. *Lexical Rules in the Hierarchical Lexicon..* PhD dissertation, Stanford University, Department of Linguistics.
- Gazdar, G. 1987. Linguistic applications of default inheritance mechanisms. In: Whitelock, P., Wood, M., Somers, H., Johnson, R., & Bennett, P. (eds.) *Linguistic theory and computer applications* (pp. 37-67). London: Academic Press.
- Hudson, R. 1984. *Word Grammar*. Oxford: Basil Blackwell.
- Jacobs, P.S. 1985. *A knowledge-based approach to language production*. Report No. UCB/CSD 86/254. Ph.D. dissertation, University of California, Berkeley.
- Keene, S. 1989. *Object-Oriented Programming in Common Lisp*. Reading, Mass.: Addison-Wesley.
- Kempen, G. 1987. A framework for incremental syntactic tree formation. In: *Proceedings of the 10th IJCAI*, Milan (pp. 655-660). Los Altos: Morgan Kaufmann.
- Pereira, & Shieber, 1984. The semantics of grammar formalisms seen as computer languages. *Proceedings of Coling*.
- Pollard, C.J. and I.A. Sag. 1987. *Information-based Syntax and Semantics. Volume 1, Fundamentals*. CSLI Lecture Notes, 13, Chicago.
- Russell, G, J. Carroll, and S. Warwick. 1990 Multiple Default Inheritance in a Unification-Based Lexicon. *Proceedings First International Workshop on Inheritance in NLP*, Tilburg.
- Steels, L. 1978. *Frame-based knowledge representation*. Working Paper 170, MIT AI Laboratory, Cambridge, MA.
- Shieber, S.M. 1986. *An introduction to unification-based approaches to grammar*. CSLI Lecture Notes 4. Stanford: CSLI.
- Touretzky, D.S., Horty, J.F. & Thomason, R.H. 1987. A clash of intuitions: The current state of nonmonotonic multiple inheritance systems. In: *Proceedings of the 10th IJCAI*, Milan. Los Altos: Morgan Kaufmann.
- Van der Linden, E., Brinkkemper, S., De Smedt, K., Van Boven, P., and Van der Linden, M. 1989. The representation of lexical objects. In: Magay, T. & Zigány, J. (eds.) *Proceedings of the EURALEX Third International Congress*, Budapest, 4-9 September 1988. Budapest: Akadémiai Kiado. (Also published as Internal Report 88-ITI-B-33, TNO, Delft).

- Van Marcke, K. 1987. KRS: An Object-Oriented Representation Language., *Revue d'Intelligence Artificielle*, 1, No. 4.
- Weinreb, D. & Moon, D. 1980. *Flavors: message passing in the Lisp Machine*. Memo AIM-602, MIT, Cambridge, MA.
- Zwicky, A. 1985. How to describe inflection. In *Proceedings of Berkeley Linguistic Society*.

SUMMARY OF ITK RESEARCH REPORTS

No	Author	Title
1	H.C. Bunt	On-line Interpretation in Speech Understanding and Dialogue Systems
2	P.A. Flach	Concept Learning from Examples Theoretical Foundations
3	O. De Troyer	RIDL*: A Tool for the Computer-Assisted Engineering of Large Databases in the Presence of Integrity Constraints
4	E. Thijsse	Something you might want to know about "wanting to know"
5	H.C. Bunt	A Model-theoretic Approach to Multi-Database Knowledge Representation
6	E.J. v.d. Linden	Lambek theorem proving and feature unification
7	H.C. Bunt	DPSG and its use in sentence generation from meaning representations
8	R. Berndsen en H. Daniels	Qualitative Economics in Prolog
9	P.A. Flach	A simple concept learner and its implementation
10	P.A. Flach	Second-order inductive learning
11	E. Thijsse	Partical logic and modal logic: a systematic survey
12	F. Dols	The Representation of Definite Description
13	R.J. Beun	The recognition of Declarative Questions in Information Dialogues
14	H.C. Bunt	Language Understanding by Computer: Developments on the Theoretical Side
15	H.C. Bunt	DIT Dynamic Interpretation in Text and dialogue
16	R. Ahn en H.P. Kolb	Discourse Representation meets Constructive Mathematics

No	Author	Title
17	G. Minnen en E.J. v.d. Linden	Algorithmen for generation in lambek theorem proving
18	H.C. Bunt	DPSG and its use in parsing
19	H.P. Kolb	Levels and Empty? Categories in a Principles and Parameters Ap- proach to Parsing
20	H.C. Bunt	Modular Incremental Modelling Be- lief and Intention
21	F. Dols en H. Daniels	Nog niet verschenen
22	F. Dols	Nog niet verschenen
23	P.A. Flach	Inductive characterisation of da- tabase relations
24	E. Thijsse H. Daniels	Definability in partial logic: the propositional part
25	H. Weigand	Modelling Documents
26	O. De Troyer	Object Oriented methods in data engineering
27	O. De Troyer	The O-O Binary Relationship Model
28	E. Thijsse	On total awareness logics
29	E. Aarts	Recognition for Acyclic Context Sensitive Grammars is NP-complete
30	P.A. Flach	The role of explanations in in- ductive learning
31	W. Daelemans, K. De Smedt en J. de Graaf	Default inheritance in an object- oriented representation of lin- guistic categories

Bibliotheek K. U. Brabant

17 000 01113217 3