

Semi-productive polysemy and sense extension

Ann Copestake

(University of Cambridge Computer Laboratory,
University of Stuttgart and CSLI)

Ted Briscoe

(University of Cambridge Computer Laboratory
and Rank Xerox Research Laboratory, Grenoble) ¹

Abstract

In this paper we discuss various aspects of systematic or conventional polysemy and their formal treatment within an implemented constraint based approach to linguistic representation. We distinguish between two classes of systematic polysemy: constructional polysemy, where a single sense assigned to a lexical entry is contextually specialised, and sense extension, which predictably relates two or more senses. Formally the first case is treated as instantiation of an underspecified lexical entry and the second by use of lexical rules. The problems of distinguishing between these two classes are discussed in detail. We illustrate how lexical rules can be used both to relate fully conventionalised senses and also applied productively to recognise novel usages and how this process can be controlled to account for semi-productivity by utilising probabilities.

1 Introduction

Discussion of polysemy has been central to much recent work on lexical semantics. Most of the arguments for (or against) attempting a fine-grained classification of semantic structure in the lexicon rest on the treatment of polysemic behaviour and attendant syntactic effects. In this paper, we argue for a distinction between two classes of systematic polysemy: constructional polysemy, where a single sense assigned to a lexical entry is contextually specialised, and sense extension, which predictably relates two or more senses. We present a unification based formalisation and implementation in which the former is treated as instantiation of an underspecified lexical entry and the latter as a rule-governed relation between signs.

It is important to distinguish putatively systematic or conventional polysemy from homonymy or un-systematic and idiosyncratic polysemy;² the two familiar senses of *bank* as ‘financial institution’ and ‘raised earth’ are homonyms, whilst the verbal sense meaning to ‘put money in a bank’ is polysemous with the nominal financial institution sense. It seems plausible that this case of polysemy is an example of a systematic sense extension by which nouns denoting artifacts become verbs denoting a purpose to which those artifacts can be put (e.g. *button*, *hammer*, *butter*, *waltz*, and so forth); though, of course, such claims need to be carefully argued for each such case.³ In what follows, we will be concerned only with cases of putatively systematic polysemy and sense extension which extend to semantically-defined classes of lexical items.

Some work on systematic polysemy has emphasised the conceptual or cognitive nature of the transfers or mappings which underlie such processes (e.g. Nunberg, 1978; 1979; Lakoff and Johnson, 1980; Fauconnier, 1985; Martin, 1990). This work is important in mapping out the range of possible conceptual transfers available and also motivating their existence. However, alone it cannot account for all aspects of the linguistic phenomena. Other work has emphasised more the conventional nature of certain transfer processes (e.g. Apresjan, 1973; Ostler and Atkins, 1992), their similarity to derivational morphological rules (e.g. Copestake and Briscoe, 1992), and cross-linguistic differences in their patterns of realisation and conventionalisation (e.g. Nunberg and Zaenen, 1992). Still further work has emphasised the intricate connection

¹We would like to thank Geoff Nunberg for many helpful comments on a draft version of this paper, much discussion and several examples. We have also benefited from the discussion at presentations of earlier versions of this paper at the Dagstuhl seminar on ‘Universals in the Lexicon’ (March, 1993) and at the CSLI workshop on ‘Ambiguity and Underrepresentation’ (September 1993). We thank two anonymous referees for their helpful comments and suggestions. We take full responsibility for any remaining errors and infelicities. This work was partly supported by the ESPRIT Acquilex-II, project BR-7315, grant to Cambridge University. We would also like to express our gratitude to Xerox PARC for providing Ann Copestake with a pleasant and productive working environment while this paper was being written.

²We use ‘conventional’ to refer to a sense which is accepted and well-attested within a speech community; sometimes this is called ‘institutionalised’ (e.g. Bauer, 1983:48) or ‘established’ (e.g. Cruse, 1986:68).

³See Clark and Clark (1979) and Hale and Keyser (1993) for two widely differing views of such denominal verbs.

between polysemy (or paradigmatic change) and associated syntagmatic effects, for example on argument structure (e.g. Levin, 1993), and the possibility of characterising some apparent polysemy as a product of syntagmatic combination (e.g. Pustejovsky, 1991, 1993).

Sense change or extension accompanies many if not most operations in the lexicon, including those familiar from derivational morphology, many grammatical function changing operations, and so forth. Some have been extensively studied, though usually more from the perspective of the morphological or syntactic consequences of such operations. In what follows, we will focus on processes of conversion or zero-derivation and particularly on processes which do not affect the major category status of the modified word. One reason for this restriction is that there is a consensus that morphological processes involving explicit affixation are rule-governed, and increasingly the focus of discussion of such examples is on their semantic effects (e.g. Riehemann, 1993); on the other hand, processes of conversion with minor or no grammatical corollaries have a more controversial status, and the need to treat these as rule-governed requires more careful argumentation. Furthermore, even if we can show that such processes can be systematic it remains to demonstrate that systematic polysemy is achieved via operations analogous to morphological rules. In this paper, we argue that processes of both sense modulation and sense change (see e.g. Cruse, 1986:50f) play a role in accounting for systematic polysemies. We attempt to distinguish modulation from change using tests traditionally associated with the distinction between vagueness and ambiguity and relate this to the formal representation.⁴

Many types of conversion process are recognised as paralleling analogous processes of derivation or compounding, and thus treated as rule governed cases of ‘zero-derivation’; for example, it is uncontroversial to suggest that a noun such as *purchase* is deverbal and ambiguous between eventive and resultative readings in the same manner as the morphologically complex *replacement*, and to propose that the lexical rule which forms deverbal nouns should cover both cases. Similarly, Hale and Keyser (1993) propose that the process of noun incorporation which forms denominal verbs in examples such as *babysit* (e.g. Baker, 1988) be generalised to account for ‘total incorporations’, that is, conversions, of the form *shelve* (from *shelf*), *calve* (from *calf*), and so forth. Likewise, Levin (1993) lists many verbal diathesis alternations which are usually treated as rule-governed conversions because of their clear effects on argument structure (e.g. causative-inchoative *He broke the glass / The glass broke*).

By contrast, apparently systematic polysemy or sense extension which at most involves subtle grammatical changes, such as various types of nominal metonymy, are often explicated in terms of processes of conceptual transfer or mapping (e.g. Lakoff, 1987), and are usually treated as essentially pragmatic phenomena (e.g. Nunberg, 1979). However, some nominal metonymies have closely-related derivational counterparts; for example, the conventional metonymy which allows a container to stand for its contents (*He drank a whole bottle (of whiskey)*) is paralleled by suffixation with *-ful* (*He drank a (?whole) bottleful (of whiskey)*).⁵ Cross-linguistically, metonymies which involve no syntactic change in English can involve systematic changes in other languages; for example, the conventional nominal metonymy by which a fruit or nut denotes the tree of the fruit or nut (e.g. *apple*, *chestnut*) is normally accompanied by a change of gender (masculine tree) in Spanish (e.g. *aceituna/aceituno* (olive) or *pomela/pomelo* (grapefruit)) and Italian (Soler and Marti, 1993). Whilst the underlying explanation for the possibility of such processes may rest on a cognitive account of conceptual transfer (Lakoff and Johnson, 1980; Lakoff, 1987) and/or a general pragmatic account of the ‘cue-validity’ of different metonymic functions (Nunberg, 1979), these cross-linguistic differences and the similarities to other rule-governed lexical processes suggest that the pragmatic account must be overlaid with an account of lexical licenses (Nunberg and Zaenen, 1992) or lexical rules (Copestake and Briscoe, 1992), in which conventionalised and language specific aspects of these general processes of conceptual transfer are expressed, and which serve as language specific ‘filters’ on the general process.

Polysemy as sense modulation through specialisation or broadening of meaning in context is intuitively a common process. Many examples that lexicographers tend to treat as alternative senses are, in principle,

⁴The term ‘vagueness’ has been used to refer to more general less specified senses, such as the ‘humankind’ sense of *man*, as opposed to the fuzzy peripheral denotation of *cup* or *game*. Cruse (1986:81) argues that ‘generality’ would be more appropriate to the former. We continue to use ‘vague’ to mean general or unspecified in deference to existing usage. The distinction between sense modulation and sense change is similar to Bierwisch’s (1982) distinction between conceptual shift and conceptual specification.

⁵The semantics of these two processes are not identical: *-ful* suffixation has an additional entailment of fullness or completeness which accounts for the preferred usage of *-ful* nominals as measure phrases (e.g. *A spoon / spoonful of sugar* in a recipe context). Such differences are expected given blocking / preemption by synonymy (e.g. Aronoff, 1976; §6 and below).

amenable to this approach; for instance, Atkins and Levin (1992) identify two senses of *reel* appropriate to the interpretation of *film reel* and *fishing reel* and demonstrate that some but not all extant conventional dictionaries list these two senses. Often the precise relationship between the premodifier and the noun is treated as a question of pragmatics (e.g. Hobbs *et al.*, 1990; Alshawi, 1992:211). However, if *reel* is defined as a container artifact with the purpose of (un)winding, where the material to be wound is left largely unspecified in the basic entry, then this definition can be specialised with the appropriate material by instantiation of the object of the (un)winding. This approach would be adequate to characterise the contribution of the premodifier to the semantics of the phrase for the two examples above. However, physical differences between types of reel would be treated as outside the domain of lexical semantics. Pustejovsky (1991) develops a theory of lexical semantics in which this approach to sense modulation can be couched. Under this account the representation of nouns includes a specification of their *qualia structure*, which encodes the form, content, agentive and telic (purpose) roles. Thus the telic role of the basic sense of *reel* would be partially instantiated. In general, Pustejovsky suggests that the notion of semantic composition be enriched to one of ‘co-composition’ in which aspects of the nominal semantic representation are integrated with aspects of the premodifier’s semantics, using a combination of type shifting of the predicate and type coercion of the nominal complement (Pustejovsky, 1993). A related phenomenon is the broadening of a sense in context; for example, *cloud* seems to have as a “mass of water vapour” basic sense, but an extended usage as a mass of anything floating *dust cloud*, *cloud of smoke*, or *cloud of mosquitoes*. One thing that normally characterises such usages is the explicit contextual specification of the way in which the sense has been broadened: thus we might treat the basic sense as taking a default content qualia value which can be overridden by a modifying phrase.

In what follows, we explore the hypothesis that systematic nominal polysemies of the kind outlined above can be divided into two types of process which we term constructional polysemy (sense modulation) and semi-productive sense extension (sense change). In constructional polysemy, the polysemy is more apparent than real, because lexically there is only one sense and it is the process of syntagmatic co-composition (Pustejovsky, 1991) which causes sense modulation. Nevertheless, we argue that the range of possible modification in co-composition is lexically specified, though pragmatically defeasible. Many cases of pre- or post- nominal modification, such as the examples of specialisation and broadening above, as well as verbal logical metonymies can be analysed in this fashion. Sense extension, on the other hand, requires lexical rules which create derived senses from basic senses, often correlating with morphological or syntactic changes. Sense extension rules are semi-productive and susceptible to processes such as blocking or preemption by synonymy, and are, we argue, formally identical to other rules of conversion and derivational morphology. Many cases of conventional nominal metonymy, such as those introduced above, can be analysed in these terms.

In §2 we describe the lexical representation language that we have developed to represent basic lexical entries and characterise systematic lexical processes. In §3 we return to constructional polysemy and motivate a more detailed analysis of specialisation as well as discussing broadening in this framework. In §4 we discuss sense extension proper with respect to grinding, portioning and other types of nominal metonymy; we address the issues of the directionality of sense extensions, their apparent ability to apply to phrases in some cases, and their productive yet highly conventionalised nature. In §5 we consider cases of ‘co-predication’ (Pustejovsky, 1994), where distinct senses are accessible for coordination and modification, and present an analysis of some cases of co-predication compatible with our accounts of constructional polysemy and sense extension. In common with other lexical processes, sense extension is semi-productive in that it is susceptible to blocking and sensitive to frequency effects; in §6 we argue that these properties can be captured by adopting a probabilistic interpretation of lexical rules and utilising probabilities in a natural fashion in language production and interpretation.

2 The Lexical Representation Language

The language we will use to represent these classes of polysemous behaviour is the lexical representation language (LRL) developed for the ACQUILEX lexical knowledge base system (LKB). The LRL is a typed feature structure language (Carpenter, 1992), augmented with defaults and lexical rules. Types are used to structure lexical entries, which are represented as feature structures (FS), and specify how they combine

by means of grammar rules, or alternatively, by constraints on phrasal types.⁶ The LRL could be used to implement a range of unification and constraint based approaches. The approach taken in this paper can be regarded (roughly) as combining an HPSG-like approach to syntax with Pustejovsky’s notion of qualia structure.

Earlier versions of the LRL have been described in Copestake (1992, 1993a,b) and we will only provide a brief sketch of the formalism here. In this paper, however, we will make use of an improved notion of default unification, which is order independent and allows for persistent defaults (Lascarides *et al.* (forthcoming), see §2.2 below). Most previous definitions of default unification have assumed that it involves incorporating into a non-default FS all the consistent information from a default FS, making no distinction in the result between information which arose from the default and non-default structures. In our treatment, by contrast, information in FSs may be marked as default (or non-default), and this distinction persists throughout subsequent default unification operations. Another difference is an improved treatment of ‘lexical’ rules, which can now operate on both lexical and phrasal signs (see §2.3). Partially specified phrasal signs can also be represented within the LRL. In general terms, we are aiming at a formalism which is adequate to represent the conventionalised, non-fully productive aspects of the language, including words, idioms and sense extension processes (which may be applicable to phrases as well as words — see §4.3). We will use *lexical* broadly to include any such specification.⁷

2.1 Types

The LRL uses a definition of typing that largely follows Carpenter (1992). The types are organised as a lattice, with top (\top) being the most general type and bottom (\perp) indicating inconsistency. This lattice, in effect, specifies compatibility between types (any two types must have a unique greatest lower bound in the lattice — they are compatible/unifiable if this is not \perp) and also allows for inheritance of constraints from types to subtypes (see Figure 1). Constraints on types are themselves FSs, which will subsume all well-formed FSs of that type — the only features that may be present on the node of a well formed FS are those appropriate to the type labelling it (see Figures 2 and 3). Furthermore, the type hierarchy itself is interpreted as constraining the class of totally specified or ‘ground’ FSs, since it is assumed to be complete, with subtypes fully covering their supertypes. That is, given t and t' are subtypes of t'' , anything of type t'' must be resolved to be either t or t' . The process of type resolution can be used to drive parsing and generation.

2.2 Lexical descriptions

In the LKB, the type language is augmented with a lexical description language that incorporates lexical rules and default inheritance. Lexical entries are defined in terms of types, for example:

```
book 1
  < > = lex-noun-sign
  < QUALIA > = art_phys
  < QUALIA TELIC PRED > = read
  < QUALIA FORM > = indiv.
```

(Here we continue to use the simple type system defined in Figure 2.) The FS is defined to have overall type **lex-noun-sign** and to have the QUALIA appropriate for an individuated physical artifact with a telic role instantiated to *read*. The ORTH feature is instantiated with a string constructed from the entry’s label, **"book"** (string types do not have to be explicitly listed in the system).

Lexical descriptions are evaluated to produce *psorts*, which are simply named FSs. We make use of psorts rather than define distinct types for each lexical entry mainly because we have found the restrictions on the type system to be inappropriate for lexical entries — we discuss this in more detail below. Various inheritance relationships are defined to operate on psorts. In theory, arbitrary parts of FSs can be related

⁶LKB and LRL are thus something of a misnomer, since the system is not specific to lexical representation, but is also used for syntagmatic description.

⁷We assume that the lexicon includes everything which is not completely compositional, that is not regularly composed from the usual meanings that the components have in isolation.

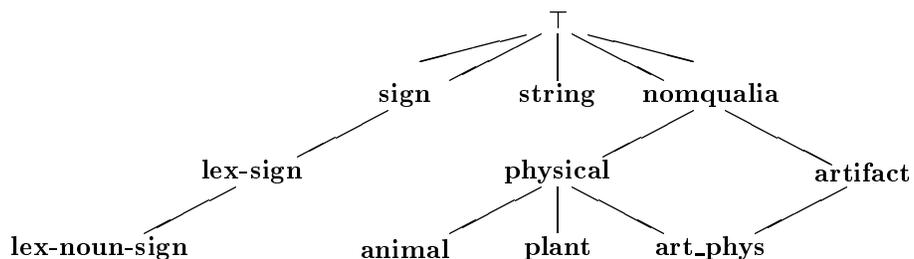


Figure 1: A fragment of a type hierarchy

top ().	gender (top) (OR male female).
string (top).	plant (physical).
sign (top) < ORTH > = string.	artifact (nomqualia) < TELIC > = verb-sem.
lex-sign (sign).	art_phys (physical artifact).
lex-noun-sign (lex-sign) < QUALIA > = nomqualia.	verb-sem (top) < IND > = eve < PRED > = string < ARG1 > = < IND > < ARG2 > = obj < ARG3 > = obj.
nomqualia (top).	sem (top).
physical (nomqualia) < FORM > = form.	eve (sem).
form (top) (OR mass indiv plural).	obj (sem).
animal (physical) < FORM > = indiv < SEX > = gender.	

Figure 2: Description of illustrative type system

$$\left[\begin{array}{l} \mathbf{art_phys} \\ \mathbf{FORM} = \mathbf{form} \\ \mathbf{TELIC} = \left[\begin{array}{l} \mathbf{verb-sem} \\ \mathbf{IND} = \boxed{\square} \mathbf{eve} \\ \mathbf{PRED} = \mathbf{string} \\ \mathbf{ARG1} = \boxed{\square} \\ \mathbf{ARG2} = \mathbf{obj} \\ \mathbf{ARG3} = \mathbf{obj} \end{array} \right] \end{array} \right]$$

Figure 3: Expanded constraint on **art_phys**

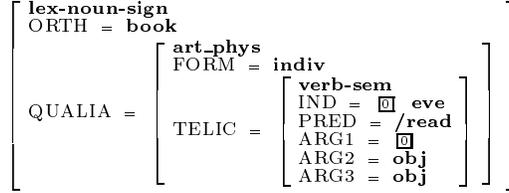


Figure 4: FS for *book*

by inheritance. In practise, we make use of two classes of inheritance specification much more extensively than others. One of these is inheritance of qualia structure, the other is used in describing a lexical entry as being derived via a productive rule, but having some exceptional value for orthography, syntax or semantics. We will concentrate on qualia inheritance here since it is more relevant to the subsequent discussion, but see Copestake (1992) for a treatment of lexical exceptions in the LKB.

We assume that the possible qualia structures can be regarded as a conceptual hierarchy (actually a lattice), certain regions of which will be associated with particular lexical entries. It is convenient to be able to describe some lexical entries as inheriting their qualia structure from others (see Copestake, 1992, 1993a). For example:

```

novel 1
  < QUALIA > < book_1 < QUALIA >.

```

states that the lexical entry for a particular sense of *novel* inherits its qualia from (a particular sense of) *book*. (The symbol < indicates inheritance.) Given this specification, *novel* would inherit its telic role from *book*. One effect of this is that it would predict that the normal interpretations of (1a) and (1b) below would both involve a reading event (see Pustejovsky, 1991 and §3, below).

- (1) a John enjoyed the book.
- b John enjoyed the novel.

However, inheritance of individual qualia must be defeasible. For example, *dictionary* should also be defined to inherit its qualia structure from *book* but has a telic role of **refer_to** rather than **read**. Default inheritance in the LKB is now formalised in terms of persistent default unification (PDU). We will give only a brief description of this here: it is fully defined in Lascarides *et al.* (forthcoming). This treatment of typed default unification is an improvement over that used previously in the LKB (Copestake, 1992,1993a) in that it is order independent and allows for persistent defaults. This allows us to define multiple orthogonal default inheritance in the lexicon in a manner which is fully declarative. Furthermore, the earlier definition of default inheritance in terms of a default unification operation applying to normal FSs, was restricted in applicability to lexical descriptions, but defaults may now persist outside the lexicon. Thus defaults may be combined during the interpretation/generation of a sentence and defaults which originate from lexical specifications can interact with pragmatic processing. In our new definition, parts of FSs may be defeasible; this is a necessary condition for default unification to be associative. In this respect, PDU is similar to the notion of defaults in Young and Rounds (1993), but their approach is limited in that their definition is restricted to non-reentrant values and in that they assume an untyped framework. In contrast, PDU uses the type hierarchy to prioritise defaults.

We use a slashed notation for partially defeasible FSs where values to the left of the slash are indefeasible and those to the right defeasible (*indefeasible/defeasible*). We abbreviate this to */defeasible* where the indefeasible value is uninteresting (e.g. where it is \top) and omit the slash when there is no (interesting) defeasible value. So, for example, the FS for *book*, shown in Figure 4, specifies that the value for the telic predicate is defeasible. The description given below for *dictionary* specifies that it inherits its qualia structure from *book* but the specific default value **refer_to** overrides the inherited value of the telic predicate.

```

dictionary 1
  < QUALIA > < book_1 < QUALIA >
  < QUALIA TELIC PRED > = /refer_to.

```

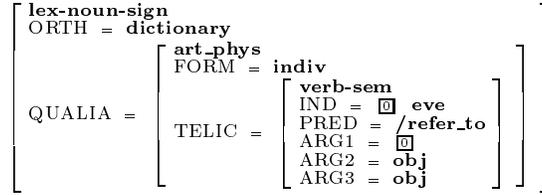


Figure 5: FS for *dictionary*

We specify the value of the telic predicate to be defeasible here as well, because for some dictionaries this might not be appropriate (e.g. Bierce’s Devil’s Dictionary) and also because the contribution of the telic role to interpretation of a particular sentence is potentially defeasible. The corresponding FS is shown in Figure 5.⁸

One effect of the difference in telic role between *book* and *dictionary* is due to the different aspectual properties of the predicates; **read** can describe a process but **refer_to** is point-like. Since *enjoy* selects for a process, (2) is odd.

(2) ? John enjoyed the dictionary.

The importance of the defeasibility of parts of the qualia structure is discussed briefly in §3 and at more length in Lascarides *et al.* (forthcoming). The persistence of the defaults ‘outside’ the lexicon is irrelevant for much of this paper, so for the most part we can continue to assume the formal account of the LKB provided by Copestake (1992; 1993b) and we, therefore, omit further discussion of PDU.

There are a number of reasons for not defining lexical entries to be types themselves. We want to maintain a distinction between the types, which are used for description or classification, and the data which they are being used to classify — i.e. the lexical entries. The type system is assumed to be complete, but we do not want to make this assumption about hierarchically arranged lexical entries. It should not be necessary or even possible to introduce features which are specific to particular lexical entries. The hierarchical organisation of the psorts is used for inheritance of information, but not for classification of words. Furthermore the condition imposed on the type hierarchy, that a unique greatest lower bound must be explicitly specified for all compatible types, is too restrictive to apply to the lexical entries, or parts of lexical entries, that we refer to as psorts. The FSs, of course, do form a lattice, but the points that are being specifically identified as psorts do not. Psorts are a way of identifying particular points in the lattice, but which points are so identified is not constrained in any way.

Furthermore, making lexical entries types obviously leads to a proliferation of types. This is particularly acute if we wish to make some lexical entries underspecified with respect to the lexical types. For example, suppose we wished to make *truth* underspecified with respect to the two types **lex-count-noun** and **lex-uncount-noun** which were both defined as subtypes of **lex-noun**. Simply specifying **truth** as an additional subtype of **lex-noun** would not achieve the correct results, since it would then not unify with a FS of type **lex-count-noun** or **lex-uncount-noun**. We would have to explicitly define **truth-count** as a subtype of **truth** and **lex-count-noun** and similarly for **truth-mass** (which means there would be no advantage of economy of representation in the underspecification).⁹ Instead we define lexical entries as FSs, but give them a special status in that they are identifiable and constrain the results of evaluating FSs which have lexical types.

In the current version of the LRL, we define psorts as constraints on certain types. If a type is defined as being *lexical* it is assumed to be constrained such that any FS to which it is resolved must be subsumed by one or more psorts of the appropriate type. For example, Figure 6 shows a FS and the possible resolutions, given the psorts shown and the types in Figure 1. If a type is defined as being *phrasal* it will normally be resolved as being constructed from lexical types, which will be constrained by lexical psorts. However it is also possible for phrasal psorts to be defined which allow an alternative analysis of the phrase. These will not

⁸We have assumed here for ease of exposition that the constraint specifications in the type system are all non-defeasible, although this will not be true in general. Type resolution, however, is determined by the indefeasible constraints and there is no notion of a ‘default link’ in the type hierarchy itself, so the formalisation of the type system itself remains very similar.

⁹In any event, there would be severe practical problems in constructing such a system, given that the type system would have to be recompiled each time a lexical entry was added.

rabbit

<> = **lex-noun-sign**
< QUALIA > = **animal**.

bull

<> = **lex-noun-sign**
< SEX > = **male**.

Query FS: $\left[\begin{array}{l} \mathbf{lex_sign} \\ \mathbf{ORTH} = \mathbf{string} \\ \mathbf{QUALIA} = \left[\begin{array}{l} \mathbf{nomqualia} \\ \mathbf{SEX} = \mathbf{male} \end{array} \right] \end{array} \right]$

Resolved FSs: $\left[\begin{array}{l} \mathbf{lex-noun-sign} \\ \mathbf{ORTH} = \mathbf{rabbit} \\ \mathbf{QUALIA} = \left[\begin{array}{l} \mathbf{animal} \\ \mathbf{FORM} = \mathbf{indiv} \\ \mathbf{SEX} = \mathbf{male} \end{array} \right] \end{array} \right], \left[\begin{array}{l} \mathbf{lex-noun-sign} \\ \mathbf{ORTH} = \mathbf{bull} \\ \mathbf{QUALIA} = \left[\begin{array}{l} \mathbf{animal} \\ \mathbf{FORM} = \mathbf{indiv} \\ \mathbf{SEX} = \mathbf{male} \end{array} \right] \end{array} \right]$

Figure 6: Constraint resolution with lexical constraints

be fully resolved FSs, but partially specified ones which will themselves be subject to constraint resolution. (This mechanism might also be used in the treatment of idioms and other (partially) fixed phrases.)

2.3 Lexical rules

Lexical rules are formalised in the LKB as feature structures of type **lexical-rule**, which has the constraint:

$$\left[\begin{array}{l} \mathbf{lexical-rule} \\ 0 = \mathbf{lex_sign} \\ 1 = \mathbf{lex_sign} \end{array} \right]$$

Application of a particular lexical rule simply involved unification of the input of the psort with the input part of the lexical rule, indicated by the path <1>, and returns the instantiated output of the rule, given by the path <0>.

An example of a lexical rule in this system is **portioning**, which covers the sense extension involved in usages such as *three beers*, where a mass noun which denotes some food or drink is converted to a count noun denoting some (conventionally served) portion of that substance. The FS in Figure 7 describes this rule using the type system from Copestake (1992) (the justification for the particular details of the representation adopted can be found there). The qualia types **c.obj** and **c.subst** indicate an edible object and substance respectively. The rule would apply to a lexical entry such as that shown for *beer* in Figure 8. Morphological rules are formally identical to sense extension rules, except in specifying a change of phonology/orthography.

One immediate question is how the notion of lexical rules fits into a constraint based framework. In Copestake and Briscoe (1992), lexical rules were essentially indistinguishable from grammar rules, and could in fact apply to phrases. This allowed us to deal with some examples of phrasal sense extension. For example, the **place** -> **group** sense extension applies both to place denoting words such as *village* and to some phrases, as in (3) (see §4.3 below, for further details).

(3) The south side of Cambridge voted Conservative.

But treating lexical rules as operating as unary grammar rules is unattractive — it obscures the distinction between the syntagmatic component of the system and the semi-productive paradigmatic component. Furthermore this treatment does not carry over in a simple way to a constraint based approach. Within a strictly constraint based framework there have been essentially three proposals for lexical rules:

1. Lexical rules expand the lexicon in a preliminary processing phase. This is the standard approach (e.g. Pollard and Sag, 1987) but is unattractive because it does not extend to analogous phrasal processes and because the lexicon is not finite.

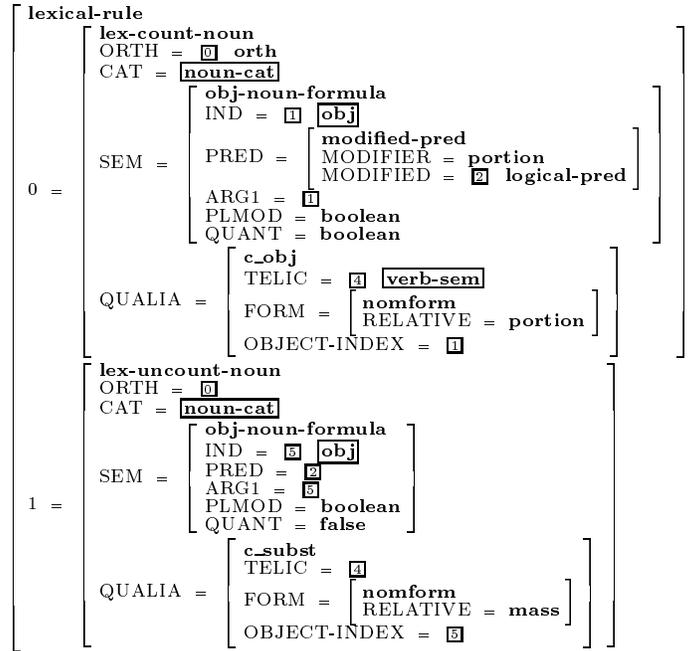


Figure 7: Lexical rule for portioning. In this figure, and subsequent examples, boxes round type labels for a node (e.g. **noun-cat**) indicate that the FS which that node heads is not shown and some features are omitted.

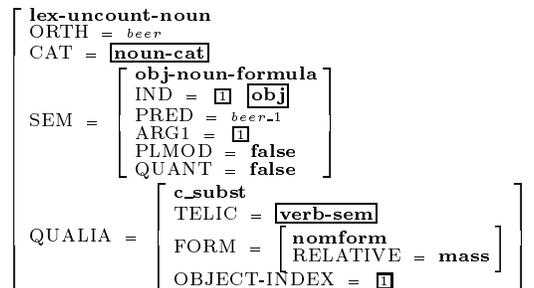


Figure 8: FS corresponding to the lexical entry *beer*

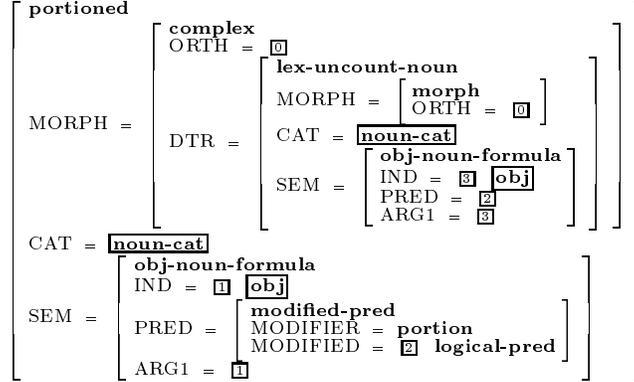


Figure 9: Portioning expressed as a complex type

2. Treat lexical rules as being similar to grammar rules, with affixes having their own lexical entries. Such an approach is suggested in Krieger and Nerbonne (1993) for derivational morphology. But for sense extension and conversion we would need to postulate zero-morphemes.
3. The place of lexical rules is taken by complex types (Riehemann 1993). For example, Figure 9, sketches a complex type which could replace the portioning rule shown before. This avoids the use of zero-morphemes for sense extension. However, it still has disadvantages — there is a proliferation of types in the hierarchy as it becomes necessary to allow lexical signs of all classes which might be formed by sense extension to be either simple or of a type that depends on their derivation. For example, **lex-count-noun** would have subtypes **simple-lex-count-noun** and **portioned**. Signs would be distinguished in this way solely because of their construction from lexical rules, which is particularly unintuitive for sense extensions since the directionality of an extension may be non-obvious (see §4.5 below). Extending the approach to phrasal signs would be possible, but would further increase the number of types. Thus this approach would work for our purposes, but the mechanics of constraint resolution are driving the representation, forcing us to postulate unnecessarily complex structures.

We treat lexical rules as generating psorts. Clearly, if we simply applied all the lexical rules to the defined psorts in a precompilation phase, this would be equivalent to the first option above. Instead of doing this, we use the lexical rules to dynamically generate alternatives during constraint resolution of nodes with lexical types. To see how this works, consider the example type system in Figure 1, but assume that instead of the type **animal** we have a type **animate**, with subtypes **animal** and **human**. Figure 10 shows a very simple lexicon and a lexical rule that converts **animal** denoting nouns to **human** denoting ones.¹⁰ The query structure shown in Figure 10 might be resolved by the lexical psort given for *grandmother*. However an alternative resolution is available via application of the lexical rule. Figure 10 shows how this is applied in effect, by ‘wrapping’ it round the query FS which instantiates the output sign of the rule and constraint resolving the result. Further resolution of the input sign, because it is matched up with a psort in the lexicon, results in specialisation of values on the output sign (the ORTH value in this case). The index [EX] is shown here to emphasise that under normal circumstances this resolution step would be part of the resolution of a sentence sign and thus the query FS shown will be part of a larger structure. Further constraints imposed on the output sign by the resolution of the surrounding structure would affect the input sign and thus limit the way in which it might be resolved. Note that this treatment implies that the output sign be resolvable with respect to the type system: it must be a potential lexical psort even though it is not actually defined as such.

This strategy involves a slight modification to the constraint resolution algorithm since it entails an external mechanism adding a node to be resolved. Resolution of this node could itself involve lexical rule application, of course, and in general, this algorithm may not terminate. This, however, also applies to the alternative formalisations. Compared with Riehemann’s approach, we are trading off greater simplicity

¹⁰We are using this as a simple example purely to explain the lexical rule mechanism, but we would, in fact, propose an animal->human rule to allow for (some aspects of) the metaphorical uses of *pig*, *worm*, *rabbit* and so on.

rabbit

<> = **lex-noun-sign**
< QUALIA > = **animal**.

grandmother

<> = **lex-noun-sign**
< QUALIA > = **human**
< SEX > = **female**.

animal-metaphor

<> = **lexical-rule**
< 0 ORTH > = < 1 ORTH >
< 0 QUALIA > = **human**
< 1 QUALIA > = **animal**
< 1 QUALIA SEX > = < 0 QUALIA SEX >.

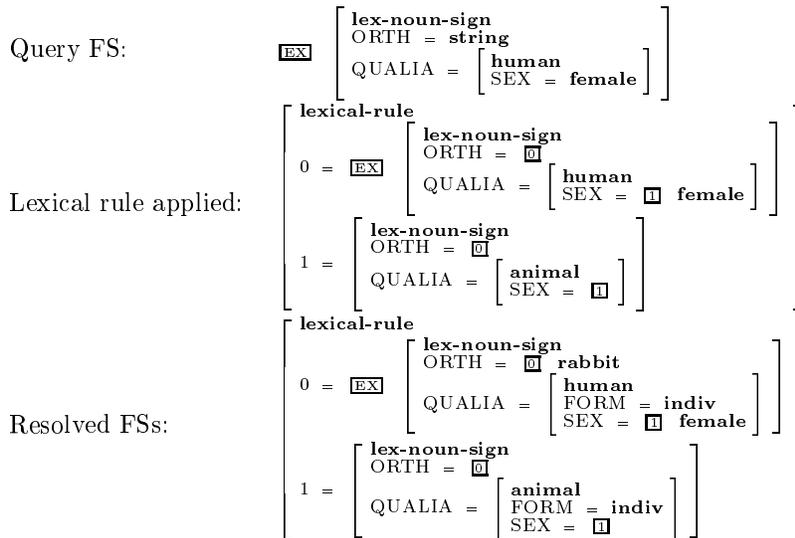


Figure 10: Constraint resolution with lexical rules

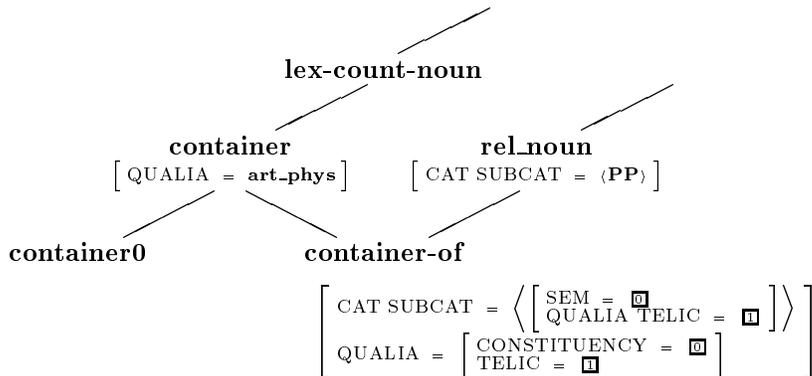


Figure 11: Outline of the description of container nouns

in the type system with a complication of the constraint resolution mechanism. From our viewpoint, one advantage is that we are maintaining a distinction between the straightforwardly syntagmatic aspects of the grammar, which are implemented by means of phrasal types, and the semi-productive processes we implement by lexical rules.

Our approach straightforwardly applies to phrases, such as example (3) above where *south side of Cambridge* denotes the group of people living there. In these cases the input form (e.g. *south side of Cambridge* denoting the place) will be a phrase with daughters (DTRS) which will themselves be further resolved in the usual way. The output structure must also be resolvable with respect to the type system. The phrasal sense extensions we have encountered so far all apply to signs which could be either lexical or phrasal as far as the context of the rest of the sentence is concerned (i.e. lexical items could be substituted for them without affecting grammaticality). Since multiword orthography does not necessitate the possession of a DTRS attribute, the lexical rule can be defined so that the output form is treated as a lexical type and will not have daughters to be resolved.

3 Constructional Polysemy

There are many cases of apparent polysemy which we would argue are better treated as ‘constructional’ polysemy, in that the lexical item is assigned one (often more abstract) sense and processes of syntagmatic combination or ‘co-composition’ (Pustejovsky, 1991) are utilised to specialise this sense appropriately. We treat this as a process of sense modulation, represented by specialisation in the LKB, in contrast to the process of sense extension to be discussed in the next section, which we represent using lexical rules.

A simple example of specialisation is the representation of *reel* in its container sense. It is reasonable to define a type **container** shown in Figure 11 that has both syntactic and semantic effects, since container nouns as a class can be subcategorised for postmodification with an *of* phrase denoting their contents (e.g. *reel of tape*) which then can be regarded as instantiating their constitutive role. Thus, the polysemy involved in the distinction between e.g. *film reel* and *fishing reel* is not regarded as lexical, and the entry for *reel* is simply:

reel 1
 <> = **container**.

The constitutive role may be instantiated by syntagmatic combination (e.g. *reel of film*) but in some cases it may only be implicit in the context.

There is, however, another source of polysemy, since container nouns as a class can also refer to their contents. Thus in a (4) *reel* can be used to refer to the film it contains.

(4) I just accidentally exposed three reels [of Ektachrome].

Furthermore, some types of polysemy will apply only to some subpart of the sense described by the lexical entry. In this particular case, *reel* used of cinema films can have an abstract sense denoting part of the film:

- (5) The mystery is only resolved in the final reel.

Here we have a sense extension from a physical object used for representation (in this case the contained object) to the abstract entity represented. Other examples of this extension will be discussed in more detail in §5.2. The point here is that it is the instantiated form of the basic entry which determines what senses are available, emphasising the need for flexible interaction between syntagmatic combination and lexical rules.

A more complex example of specialisation by constructional polysemy is adjectival premodification; it is well known that in examples such as (6), the adjectives take on different meanings depending on the nature of the modified head.

- (6) a a sad poem / poet / day
b a fast motorway / car / driver

Such examples have been used to argue that adjectives should be treated as higher-order predicates or should introduce an unspecified predicate representing the relation between the property denoted by the adjective and that denoted by the head noun (e.g. Hobbs *et al.*, 1990). Pustejovsky (1991; 1993) argues that some such adjectives can be analysed as predicates which coerce the type of the head and operate on its qualia structure. Thus he analyses *fast* as a predicate which selects the eventive qualia accessible through the entries for the head nouns in (6b).¹¹ The claim is that nouns denoting artifacts make available as part of their lexical specification an agentive and telic role representing their (typical) process of creation and of use, respectively. Similarly deverbal nouns make their underlying verbal predicate accessible in the same manner. Thus, an adjective selecting an eventive argument ‘coerces’ the type of the noun into one of the eventive qualia or the predicate underlying the deverbal noun.

Pustejovsky (1991, 1993) also discusses other examples of ‘logical metonymy’, in which the semantics of a verbal predicate and the type of its complement exhibit mismatches, such as (7).

- (7) a Sam enjoyed (drinking) the beer
b Sam enjoyed (watching) the film
c Sam enjoyed (reading) the book
d Sam enjoyed (eating) the caviar

enjoy subcategorises for a NP or progressive VP complement syntactically, but semantically requires a complement with an eventive interpretation in which the experiencer subject of *enjoy* participates as understood subject. Each of the examples in (7) is grammatical with or without the bracketed progressive participle. However, in the case where it is not present the interpretation remains (by default) identical. Analogously, to the case of adjectival modification, Pustejovsky (1993) captures the similarity between the two subcategorisation possibilities for *enjoy* by means of a type shifting operator applied to the predicate, and uses a type coercion operator which selects from the eventive qualia of the NP artifact-denoting complement to express the ‘co-compositional’ aspect of the resultant interpretation.

Briscoe *et al.* (1990) presents an analysis of logical metonymies with *enjoy* which is based on treating type coercion as a (unary) grammatical rule which alters the type and interpretation of the NP. However, Copestake and Briscoe (1992) and Godard and Jayez (1993) point out problems with this analysis stemming from possibilities of ‘co-predication’; for example, it seems quite possible to coordinate predicates which require physical objects and events as complements, as in (8).

- (8) a Sam picked up and finished his beer
b Sam ate and enjoyed the caviar
c Sam wrote but later regretted that article

¹¹Briscoe *et al.* (1990) and Godard and Jayez (1993) point out that there are problems with Pustejovsky’s technical approach to type coercion relating to co-predication (see §5 and below). We omit details of this proposal here, which is described most fully in Pustejovsky (1993).

Therefore, we treat this type of polysemy as a question of selecting the appropriate aspect of the meaning of the complement, rather than a change in the meaning of the NP itself. Traditionally, this is closest to saying that nouns denoting artifacts are vague, rather than ambiguous, between eventive and objective readings, in these contexts.

Consider first the example of *fast typist*. The effect we want is for *fast* to apply to events of the typist typing — i.e. the paraphrase of *fast typist* is (by default) *typist who types fast*. We will assume that we do this by reifying the event, giving a logical form equivalent to:¹²

$$[x][\textit{typist}(x) \wedge \textit{fast}(e) \wedge \textit{type}(e, x)]$$

We achieve this result by assuming that the qualia structure for *typist* has its telic role instantiated to:

$$[x][\textit{type}(e, x)]$$

where x is coindexed with the ‘normal’ variable.¹³ Thus the lexical entry for *typist* contains structures equivalent to the following:

$$\begin{aligned} \langle \text{SEM} \rangle [x][\textit{typist}(x)] \\ \langle \text{QUALIA TELIC} \rangle [x][\textit{type}(e, x)] \end{aligned}$$

The type **adjective** has subtypes for adjectives that select the telic role, the agentive role and so on. The basic type **adjective** is subcategorised for nouns, and has the following semantics:

$$[x][\textit{adj-pred}(w) \wedge P(w, x)]$$

The treatment is similar to that proposed by Hobbs *et al.* (1990), for example; rather than directly equating the entities denoted by the noun and the adjective, the relationship between the two, denoted above by P , is underspecified. However, in our approach, information from the qualia structure provides the instantiation. In the case of **telic-adjectives**, P will be instantiated by the telic predicate.

The lexical entry for *fast* can be specified as **adjective** with the semantics instantiated so that it can only be true of an event. Any particular instance of *fast* in an utterance will have to become resolved to one particular subtype of **adjective**. In the case of *fast typist*, the normal form of the adjective is ruled out since *typist* is object denoting and only the telic role specifies a possible predicate. The choice of predicate may be determined by selectional restrictions, which can be encoded in the LKB as constraints on the types governing the predicate argument structure, but we will not discuss the details here. The qualia structure of the modified phrase is equal to that of the noun — see Figure 12.

In this formulation the qualia structure of the noun is not itself directly modified by the adjective. This differs from the treatment we gave in Briscoe *et al.* (1990) where, because we unified the entire telic role into the representation of the modified nominal, all telic events were, in effect, modified by the adjective. This meant, for example, that the interpretation of *enjoy the long book* entailed that the reading event assumed

¹²We use a linearised form equivalent to the FS representation here for readability. We will leave some aspects of the representation incomplete where the details are not relevant to our main concerns. For example, we do not specify here how the event variable e should be bound. Simple existential quantification looks unsatisfactory since there seems to be something generic or habitual about *fast typist*. One possible approach might be to treat the domain of events as having a lattice structure (e.g. Krifka, 1987) which would allow us to make the event referred to as *fast* the composite of subevents of the typist typing (cf Ojeda, 1993, on generic nominals) or perhaps the composite of some contextually salient subevents. Since *fast* need not be fully distributive, this would not imply that all subevents were *fast*. But we have not worked out the details of such a treatment since it is not at all obvious how many typing events *fast* ought to apply to. Most work on generics and habituals makes the assumption that they can be paraphrased using *normally* or *usually*, but it is not clear that this is true of *fast typist*, *fast car* etc. It is possible to assert that Bill is a *fast typist* even if he usually types at 20 words per minute but was observed doing 120 wpm in a competition. An individual car can perhaps be truthfully said to be fast even if it has never been driven above 40mph yet, as long as its potential is known. This situation is not peculiar to this class of adjectival modification: *John eats snails*, for example, can be true even if he has only done so once or twice (cf Pelletier and Schubert’s (1988) comments on *Frenchmen eat horsemeat* and similar examples).

¹³The status of qualia structure in our approach is slightly different to that of Pustejovsky and Boguraev (1993) in that we include qualia structure in the *lexical* representation of the noun (as a component of a FS in the LKB) and specify type coercion in unification based terms. However, we also recognise the need for interaction between qualia structure derived stereotypical eventive readings and other pragmatically or contextually determined interpretations. In our account, the stereotypical reading is specified by default as a by-product of the parsing process, but can be overridden pragmatically (see Briscoe *et al.*, 1990; Lascarides *et al.*, forthcoming).

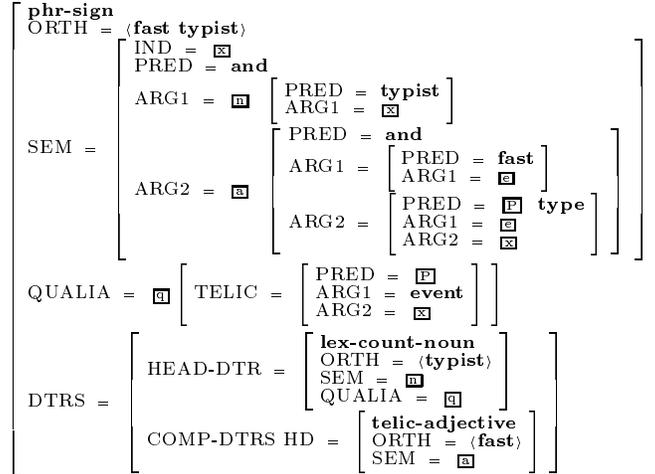


Figure 12: FS for *fast typist* (letters are used here to indicate reentrancy rather than the usual numbers to make the figure easier to follow).

was also *long*, which is not necessarily correct. In our current treatment, the variable is specified by the adjective alone and this problem does not arise.

The interpretation of *fast typist* as someone who types fast is defeasible. In the context of a race between typists and accountants, for example, a fast typist might be one who can run, ski or ride a motorbike quickly; in this case the predicate is given contextually. Briscoe *et al.* (1990) argues for the notion of a default lexical interpretation, which can be overridden in informationally rich contexts. Lascarides *et al.* (forthcoming) describes how persistent default feature structures can be used to formalise this, by specifying the portion of the semantic representation derived from the qualia structure as default.

Our current treatment of *enjoy* is similar to that of *fast*, in that the ‘coercion’ is internal to the verb semantics. (Godard and Jayez (1993) also adopt such an approach.) We treat *enjoy* as having a type which can either be specialised to take an event denoting complement in the usual way, or to introduce an indirect relationship between an object and the event, which will be instantiated via the telic role – see Figures 13 and 14.¹⁴

One further example of an operation which can be involved in constructional polysemy could be called *broadening* since usages are available in context which appear to semantically subsume the basic sense. Usually it appears that a quale which is specified in the basic sense becomes overridden in context. For example, the normal usages of *bank* and *cloud* could be specified as stating both form and composition (earth / water vapour). However, both have usages where alternative compositions are stated *bank of rhododendrons*, *bank of clouds/cloud bank*, *cloud of mosquitoes*, *dust cloud*. In some comparable cases the broadened sense may appear more metaphorical, for example *forest of hands*. In many cases there is evidence that broadening of meaning has taken place diachronically and that the original senses tended to be specific and concrete (see Sweetser 1990). It seems appropriate to regard these examples as being comparable to those given above in that there is a modulation of sense rather than a complete shift, but unlike the cases discussed above, this modulation is most naturally expressed as being non-monotonic. For example, in contrast with the case of *reel* given earlier, there is a very strong preference for one particular sense and the alternative interpretations are not conventionalised, but given by context (there is no conventional interpretation of *cloud* as *cloud of mosquitoes*). This implies that non-default interpretations will only be usual in contexts which explicitly give the exceptional component (normally by compounding or post-modification). This then, is rather similar to the situation with respect to the stereotypical readings of *enjoy the book* and similar examples (Briscoe *et al.*, 1990).

To represent broadening we make use of lexically specified persistent default components of the qualia structure and allow these to be overridden. In the FS for the lexical entry for *cloud* shown in Figure 15 the qualia structure is stated to refer necessarily to an individuated physical object of amorphous form, with a

¹⁴We leave the treatment of both *fast* and *enjoy* with respect to coordination to §5 below.

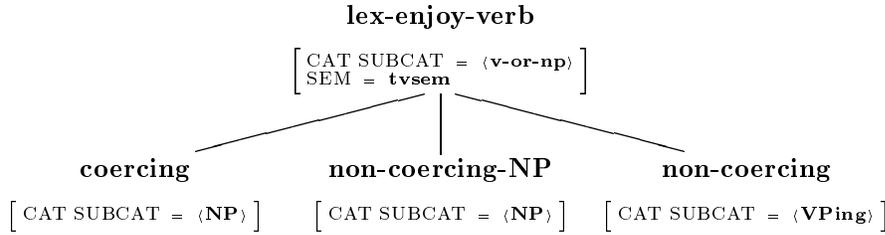


Figure 13: Outline of type hierarchy for *enjoy* and similar verbs

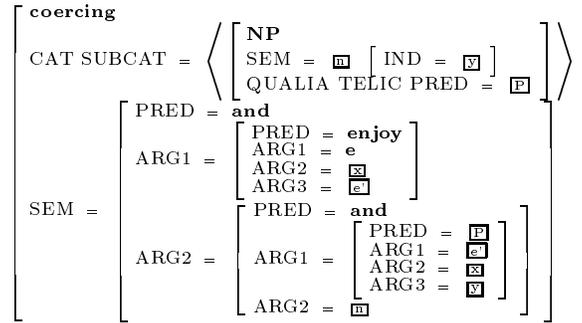


Figure 14: Coercing form of *enjoy*

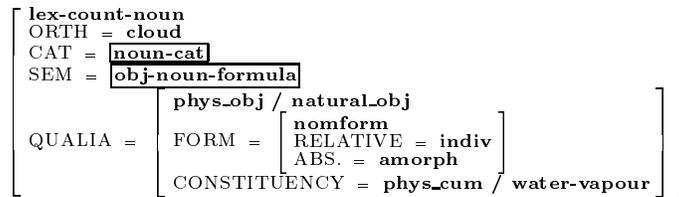


Figure 15: Lexical entry for *cloud*

composition that is also physical and refers cumulatively (i.e. the composition is either a mass or a plural object). By default, *cloud* is a natural object (as opposed to an artifact) and is composed of water vapour.¹⁵ Referring to the process of overriding the lexically specified defaults as broadening is perhaps somewhat misleading, since a more general FS never actually exists in isolation according to this treatment. The intuition that the sense is broadened is reflected in the non-defeasible components of the modified structure, however: for example the semantic contribution of *cloud* to *cloud of mosquitoes* could be represented as a FS with unspecified composition.

Broadening could alternatively be represented using a lexical rule which removes part of the qualia structure. But this is a less attractive account since it would be difficult to avoid spurious ambiguity which would occur if the broadened sense were specialised to have a structure equivalent to the usual sense. Furthermore, the default account gives a natural explanation for the fact that explicit contextual specification of the alternative compositions is necessary for the usage to be interpreted in its broadened sense, which the lexical rule account would fail to capture, without some additional mechanism. In general, we see the use of lexical rules as appropriate when there is a shift in syntactic or semantic type, as will be illustrated in more detail in the next section.

¹⁵This description has been somewhat simplified but in any case we would not claim that it is completely adequate. It does not, for instance, cover the mass use of *cloud*, found in (9a), which seems to be available only with the default usage (compare (9b)):

- (9) a We flew into dense cloud.
b * We walked into dense cloud of smoke.

Nor does it cover the metaphorical uses, such as *cloud of suspicion*.

4 Sense extensions

By contrast with constructional polysemy, we argue that there are systematic polysemies which are best represented as lexical rules, which we refer to as sense extensions; that is predictable creation of different but related senses. As described in §2, the formalism that we utilise to express these rules is equally applicable to derivational processes, as well as those of conversion, in that we treat all such lexical processes as mappings between lexical (and occasionally phrasal) signs.¹⁶ From our perspective, it is accidental that some rules specify phonological modifications whilst others do not.¹⁷ However, we concentrate on cases which involve little if any grammatical change, since these constitute the major challenge to a uniform theory of lexical processes.

The examples of sense extension discussed below could be broadly characterised as metonymic. In Briscoe and Copestake (1991), we suggested that similar mechanisms could be used to account for metaphoric processes as well. For example, the sense extension from animals into metaphorical senses denoting humans with some particular characteristic is apparently productive (e.g. *John is a lamb / pig / wombat*), although the actual characteristics involved cannot be predicted from knowledge of the animal sense. We would argue, for example, that the properties ascribed to a person by *pig* are stereotypical associations with the animal, which would not be encoded in the qualia structure. Despite the more associative or analogical nature of metaphorical sense extension, there is a core component to such processes which should be expressed in terms of a sense extension rule. In general, we assume that the possible mappings defined by sense extension rules define the limits to the possible shifts in meaning, but more general reasoning may be involved in determining the meaning more exactly in a particular context. However, in this paper, we will concentrate on metonymic examples.

4.1 Grinding and portioning

One process of sense extension is that which creates mass nouns denoting an unindividuated substance from count nouns denoting an individuated physical object of some kind. Given the right context, this process can apply quite generally. The context normally suggested is to imagine a large grinding machine, the Universal Grinder (see, e.g. Pelletier and Schubert, 1989), which would, for example, turn *a table* into some substance that could be referred to by the mass term *table*. Conventional subcases of grinding exist, for example, food-denoting mass nouns can be formed from animal-denoting count nouns (e.g. *lamb, rabbit, haddock, chicken*). This extension appears to be productive, at least in a sufficiently marked context; for example, in the LOB corpus (10) we find the use of *mole* as a mass term.

(10) Badger hams are a delicacy in China while mole is eaten in many parts of Africa.

We therefore cannot assume that the extended senses are listed explicitly in the lexicon. As in this example, where the animal sense is a count noun and the meat sense is mass, sense extensions may affect syntactic behaviour. However, the syntactic difference is not criterial since in examples such as (11) it is the predicate rather than the complement which indicates that grinding has occurred.

(11) Sam enjoyed the lamb.¹⁸

Furthermore, unlike the case of co-predication with constructional polysemy, it seems much harder to coordinate predicates selecting for the ground and unground senses of a complement, especially if this is combined with co-composition, as (12) illustrates.

(12) a ?Sam fed and carved the lamb
b ??Sam fed and enjoyed the lamb

¹⁶This makes our approach closest to that of word-based morphology (e.g. Aronoff, 1976) but with the possibility of phrasal based operations as well.

¹⁷In fact, there is more to be said on this topic, since it seems plausible that derivational rules are less ambiguous, because of the information about the process conveyed by the affix, and therefore, perhaps more fine-grained in the sense modifications they produce. Discussion of such differences though would take us outside the scope of this paper.

¹⁸Note that in (11) both grinding and co-composition are required – we assume that grinding of animals to meat creates an artifact which is specified for eventive telic and agentive qualia, leading to a default ‘Sam enjoyed eating the lamb’ interpretation.

In §5 we return to similar more acceptable cases and argue that in restricted cases such examples are comprehensible as instances of co-composition with the *ORIGIN* specification of the ground predicate. However, for the moment we assume that such examples suggest that we have a genuine ambiguity, as opposed to vagueness: in this case between animal and ‘animal stuff’ denoting senses.

One striking similarity between conventionalised cases of grinding and derivational processes is that both can be blocked (e.g. Aronoff, 1976); that is, undergo preemption by synonymy or lexical form. For example, Aronoff notes the pattern in (13) and argues that *gloriosity* is blocked by *glory*, whilst *curiosity* and *curiousness* co-exist because they are not synonymous (possibly as a result of semantic specialisation).

- (13) a curious / curiosity / curiousness
 b glorious / *gloriosity / gloriousness
 c His curiosity was attracted to the curiousness of the phenomenon
 d ??His curiousness was attracted to the curiosity of the phenomenon

Thus (13c) and (13d) are not equally acceptable because *curiousness* is typically predicated of things, unlike *curiosity* which seems more appropriate to people. Similarly, we find the examples in (14) with the conventionalised subcase of meat grinding are odd.

- (14) a ?Sam ate pig (pork)
 b ?Sam likes cow (beef)
 c ‘Hot sausages, two for a dollar, made of genuine pig, why not buy one for the lady?’
 ‘Don’t you mean pork, sir?’ said Carrot warily, eyeing the glistening tubes.
 ‘Manner of speaking, manner of speaking,’ said Throat quickly. ‘Certainly your actual pig products. Genuine pig.’
 (Terry Pratchett, 1989. *Guards, Guards!*, Gollanz, London. (p. 155, Corgi edition, 1990))
 d There were five thousand extremely loud people on the floor eager to tear into roast cow with both hands and wash it down with bourbon whiskey. (Tom Wolfe, 1979. *The Right Stuff*, Farrar, Straus and Giroux, New York (p. 298, Picador edition, 1991))

Nevertheless, such examples do occur and when they do, as in (14c,d) the intuition is that they are not synonymous with the underived senses of *pork* and *beef*; they either convey a negative attitude to the consumption of the meat on the part of the speaker or an entailment of extended denotation, where more of the cow or pig than is normally considered ‘meat’ is being treated as food. Blocking appears to be explicable on the basis of Gricean principles, in particular the Maxim of Manner. Given a choice between ways of expressing the same meaning, the most easily interpretable ones should be preferred. In general, this implies that common terms should be used rather than obscure ones, briefer/simpler forms rather than more complex ones, and unambiguous expressions instead of ambiguous ones.¹⁹ Apparent violation of this maxim carries the (discourse) implication that the terms are not strictly synonymous, thus terms which are normally blocked will be interpreted as carrying additional entailments (see Briscoe *et al.*, 1994 for additional discussion).

Nunberg and Zaenen (1992) point out that conventionalised subcases of grinding vary cross-linguistically and that there are no clear pragmatic explanations either for this variation or the absence of some conventionalised cases in English. For example, they report that in Eskimo (at least conventionalised) grinding of animals is ungrammatical; and in English it seems that grinding of fruits or nuts to produce liquids is not conventionalised: thus, the examples in (15) are awkward, though (15b) is imaginable, for example, in the

¹⁹Avoidance of ambiguity might apply to sense extension, but not to derivation and it is not obvious how to measure brevity/complexity. In fact, blocking is explicable simply in terms of avoiding obscurity, by which we mean that the speaker will generally use the form which has highest frequency. At first sight it might seem that this is circular, but note that we are not trying to account for the distribution of the blocked form in the general speech community here, but only for the effects on the individual speaker. Obviously the choices of individual speakers affect overall frequencies, giving a positive feedback effect in this case. We consider this in more detail in §6.

context of a conversation between professional cooks.

- (15) a ?I drink pear rather than peach (cf. I drink orange for breakfast)
b ?I fry courgettes with olive rather than safflower

For these reasons, they argue that a language specific system of ‘lexical licenses’ must be provided in order to specify which subcases of the more general conceptual grinding transfer occur conventionally in a particular language. In addition, different languages choose different grammatical means to encode grinding and its subcases; for instance, in Dutch meat grinding of animals is usually realised by explicit compounding of *vlees* so lamb meat is *lamsvlees* and so forth. The conversion process appears to be restricted to the more stereotypical animals which are farmed for meat, such as chicken. In this way, Dutch appears to somewhat mirror the situation in English with liquid grinding, where certain stereotypical ‘juicy’ fruit denoting nouns, such as *orange* can acquire a juice sense through grinding, but the majority require explicit compounding (e.g. *apricot juice*).

Nunberg and Zaenen (1992) also argue that the meaning of ground nouns is defeasible and therefore pragmatically specified. Thus, in the case of grinding of animals, they would provide a lexical license specifying that this is conventional in English, but argue that the interpretation of ground animal denoting nouns as meat is contextually specified. Thus, in (14a,b) the Maxim of Manner requires that we choose *pork* or *beef* because these terms have a more restricted denotation than ‘animal stuff’. On the other hand in examples such as (16a,b) the context tells us that a more restricted ‘meat’ denotation is appropriate. Whilst, in (16c) the context tells us that a ‘fur’ reading is more appropriate, and in (16d) that nothing more specific than ‘stuff’ is entailed.

- (16) a Sam eats rabbit regularly
b Sam enjoyed the rabbit
c Sam wears rabbit regularly
d Sam both wears and eats rabbit

Our approach is similar in that we posit a general abstract lexical rule of grinding and conventionalised subcases, including animal meat grinding and animal fur grinding. However, we also suggest that whilst the more specific conventionalised ‘meat’ and ‘fur’ senses are defeasible in appropriate contexts (because the more general ground sense is also available), they are specified lexically as a component of the conventionalised subcases of the grinding lexical rule.

The general rule of **grinding** is shown in Figure 16 (using the type system described in Copestake, 1992, which also discusses the formal semantic properties of the grinding function in the context of the general treatment of mass terms proposed by Krifka (1987)). The effect of the lexical rule is to create from a count noun with the qualia properties appropriate to an individuated physical object, a mass noun with properties appropriate for an unindividuated substance.

We specialise the grinding rule to allow for cases such as the animal/meat extension explicitly. The typed framework provides us with a natural method of characterising the subparts of the lexicon to which such rules should apply. The lexical rules can, in effect, be parametrised by inheritance in the type system. For example, we can give rules which inherit information from **grinding** such as **meat-grinding**:

```
meat-grinding
  < > < grinding < >
  < 1 QUALIA > = animal
  < 0 QUALIA > = c_subst.
```

As in §2.3 **c_subst** is a type which stands for normally comestible naturally derived substances. The lexical rule can be applied to the lexical entry for *rabbit* to generate a sense corresponding to ‘edible stuff derived from rabbits’ partially represented as shown in Figure 17. Here the specification of the value for the telic role arises from the constraint on the type **c_subst**. Using the notion of persistent defaults described in Lascarides *et al* (forthcoming), we can treat this as defeasible. The **meat-grinding** rule creates a second extended sense for the mass noun *rabbit* (and other animal denoting count nouns) but does not result in the full specification of what might usually be taken as the meaning of the meat/flesh sense. The substance is stated to be edible (to be precise, to have the normal purpose of being eaten) and to be derived from the

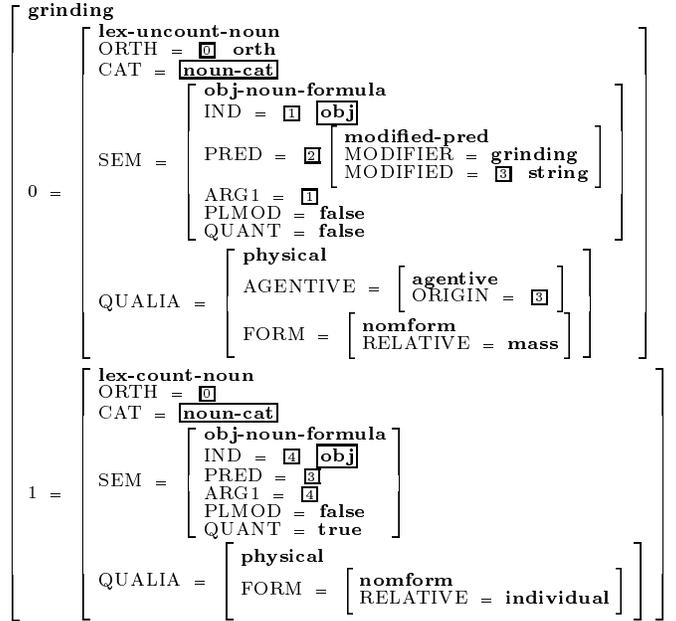


Figure 16: Grinding

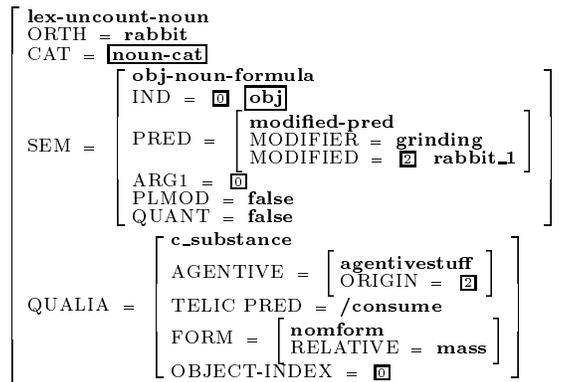


Figure 17: Meat/flesh sense of *rabbit*

animal, but there is no attempt at defining the meaning to exclude, say, stuff derived from bones; particular cultural assumptions will affect exactly what is taken to be edible, so *rabbit* will usually exclude the bones but *whitebait* will not, for example. Thus not all the characteristics are captured by the lexical rule and we assume that pragmatic effects will ensure further contextual specialisation.

The more specific rules which inherit from the general grinding rule, express the conventionalised processes that apply to semantically specified parts of the lexicon. In addition to **meat-grinding** we could also define a lexical rule which gives the fur/skin sense, available for *rabbit*, *mink*, *beaver*, *calf*, *lizard*, *crocodile* and so forth. In this way we account for the possibility of multiple distinct mass senses being possible. In context, a general mass sense corresponding to the application of the underspecified grinding rule is available, as in (17).

- (17) After several lorries had run over the body, there was rabbit splattered all over the road.

Thus, under this account, the defeasibility of the more specific sense is predicted in terms of ambiguity. The alternative of relying on pragmatic specification of a single underspecified sense seems to us less satisfactory because of the specificity of readings found in uninformative contexts; for example, in examples such as (16b) or (18), the natural interpretation is that the rabbit was eaten.

- (18) Sam enjoyed but later regretted the rabbit

Under the co-compositional account of such constructional polysemy (see §3) this is straightforward since the *meat-grinding* sense of *rabbit* provides a telic role which allows the eating interpretation to be constructed.²⁰ However, if the lexicon does not propose such a sense, it is unclear what it is about the context which allows pragmatic specialisation of the interpretation. Briscoe *et al.* (1990) provide empirical support for the hypothesis that the lexicon proposes and pragmatics disposes of such initial interpretations: on the assumption that logical metonymy will be utilised when a reading based on qualia is appropriate, or when the context is rich enough to provide determinate information to override this ‘default’; and that an explicit event will be specified where a non-default reading is appropriate, but the general context is not rich enough to override the default. Thus, a verb like *enjoy* occurs mostly with metonymic NP complements, but when it does occur with progressive VPs the interpretation is never that which would be predicted by co-composition with eventive qualia; whilst with metonymic NP complements, where the default reading is inappropriate the context is always informationally rich and determinate.

Multiple sense extensions / lexical rules may be applied in sequence. For example, we mentioned in §2.3 the lexical rule **portioning** which converts food or drink denoting mass nouns into count nouns denoting a portion of that substance (e.g. *three beers*). This is clearly productive, it can be used with names of particular types of beer, for instance, such as *three Heinekens/IPAs/Anchor Steams*. It can also apply to extended senses such as *three lambs*, at least in the context of a restaurant. This ‘feeding’ of lexical rules raises the issue of why ground portioned nouns are not, for instance, reground creating an infinite sequence of more and more derived senses. There are several potential solutions to this problem; one might be to set up the rules so that grinding feeds portioning but not vice-versa. However, we do not think that this is necessary, and in fact there is no reason to believe that portioned count nouns are of a type inaccessible to grinding. Rather we think that the non-existence of ground portioned nouns follows from the semi-productivity of lexical rules; the ground portioned sense is synonymous with the original mass sense and is thus blocked. We return to the issues of semi-productivity and blocking in §6.

4.2 Nominal Metonymies

Grinding can be characterised as a set of metonymic sense extensions in which the animal comes to stand for something derived from the animal. However, it appears to have a different flavour to many of the nominal metonymies identified by Nunberg (1979), for example. Many of these involve objects standing for people, as in (19).

²⁰We defer to §6 an explanation of why this reading is preferred to one in which Sam is wearing rabbit fur.

- (19) a The third violin is playing badly
 b The Armani suit lounging gracefully at the bar looks bored
 c London said that a new passport could not be issued
 d The village voted conservative at the last election

Although these putative sense extensions seem to have no grammatical effects, sometimes they can affect agreement. Nunberg (1979) and Pollard and Sag (in press) discuss the use of food to denote people, which is a less conventionalised example of a similar metonymy, as in (20).

- (20) a The ham sandwich wants a coke
 b The french fries is getting impatient

It is clear that agreement in (20b) is determined by the referent rather than the syntax of the NP *french fries* which would induce plural agreement given a non-metonymic reading. Similarly, co-predication of such examples seems awkward, as in (21).

- (21) a ??The ham sandwich wants a coke and has gone stale
 b ??The french fries is getting impatient and are getting cold
 c ??The third violin is scratched and playing badly
 d ??The Armani suit is at the bar and crumpled

Similarly, it is clear that pronominal agreement and reflexivisation are also affected by transfer of reference (Fauconnier, 1985; Nunberg, 1993; Pollard and Sag, in press). These observations suggest to us that these nominal metonymies must have a non-pragmatic component and must be treated as distinct senses / signs. Within our framework, we propose to treat them as sense extensions and provide lexical rules for them, analogous to those developed for grinding and portioning.

Another such sense extension is that from a word denoting a fruit (or nut) to a plant bearing that type of fruit (e.g. *apple*, *gooseberry*, *walnut*) which is found in Italian and Spanish as well as English.²¹ However, in the Romance languages the fruit is usually (but not always) feminine while the tree is masculine (there are one or two exceptions). For example, in Spanish we have *aceituna/aceituno* (olive), *pomelo/pomelo* (grapefruit) (see Soler and Marti, 1993). In a few cases, the suffix *ero* applies – *albaricoque*, *albaricoquero* (again illustrating the similarity of sense extension, conversion and derivation). The basic type for the lexical rule can be stated as:

fruit-to-tree (lexical-rule)
 < 1 > = **lex-count-noun**
 < 0 > = **lex-count-noun**
 < 1 QUALIA > = **c_nat_obj**
 < 0 QUALIA > = **plant**.

The normal lexical rule for Spanish can then be stated as:

fruit-to-tree-ESP
 <> = **fruit-to-tree**
 < 0 SEM IND AGR GENDER > = **masc**
 < 1 QUALIA AGENTIVE ORIGIN > = < 0 SEM PRED > .

The exceptional cases can be stated using explicit lexical entries which override the usual results of lexical rule application:

higuera
 <> < (*higo* + **fruit-to-tree-ESP**) <>
 < SEM IND AGR GENDER > = **fem**.

This example illustrates that some nominal metonymies, just like grinding, can have different grammatical encodings in different languages and this supports our contention that such processes should be treated as

²¹Some techniques for exploiting parallelism between lexical processes in machine translation are described in Copestake and Sanfilippo (1993).

language specific lexical rules, creating lexical entries (signs) with extended senses and different grammatical and/or phonological specifications, as required. We return to the issue of how to distinguish such cases from those of sense modulation or constructional polysemy in §5.

4.3 Phrasal sense extension

There are some examples where sense extensions apparently apply to phrases. Thus the place → group sense extension applies both to words such as *village* and place denoting phrases, as in (22).

- (22) a The south side of Cambridge voted Conservative
 b Three villages / three villages south of the river / ?three villages built of stone
 voted for the proposed ban on timber production.

These seem quite restricted; in this particular sense extension it appears that only modifiers which might apply to the group of people, or which are locational (as in *the south side of Cambridge*) are fully acceptable. With grinding too, there are cases of phrases, or at least compounds, undergoing the sense extension, as in (23).

- (23) Here you can eat alligator tail, elk, rattlesnake and that snicker-inspiring delicacy,
 Rocky Mountain oysters. (CSAA magazine)

The treatment of such phrasal sense extensions in the LKB is a straightforward generalisation of the lexical case since as we described in §2.3 ‘lexical’ rules can apply to any feature structure representing a lexical or phrasal sign with the appropriate properties.

Some examples where a sense extension apparently applies to a phrase are misleading though, since the availability of qualia structure does allow for modifiers which apply to the unextended sense. For example, in the meat grinding cases, we get *corn-fed chicken* and *young lamb*, where the adjectival phrase, on semantic grounds, has to apply to the animal, not the meat, but we also get, for example, *young veal*, *corn-fed beef*, so such examples do not demonstrate that grinding is applying to a phrase. We would analyse all these cases as ones in which the modifier is applying to the ORIGIN feature of the qualia structure (see Figure 17 and the example of *fast typist* shown in Figure 12, and also §5).²²

4.4 Novel sense extensions

Pragmatic factors clearly affect the acceptability of the underspecified, unconventionalised uses of sense extension typified by the ‘ham sandwich’ example in (22a). Something like Nunberg’s (1979) conditions on transfer of reference are needed for the intended referent to be identifiable. But these in themselves do not sufficiently delimit the possible uses of even the novel sense extensions. Nunberg postulates a set of basic transfer functions — we would identify these with our most general sense extension rules. The existence of a (unidirectional) object → human basic transfer function allows for the ham sandwich sentences, in appropriate contexts, but the converse case does not seem to be possible. Thus, for example, (24) is an unacceptable way of referring to the food that has been ordered by an identified customer.

- (24) * The man with the brown suit is in the microwave.

Nunberg discusses the cue-validity of such putative transfer functions and argues that those which occur are motivated by the value of the function as a determinant of the referent. However, *a priori* there is no apparent reason why the function from human → object cannot apply in contexts in which (24) might be uttered.

For the ham sandwich examples the basic sense extension rule that applies could be characterised as physical object → human. It seems reasonable to assume that such a rule is analogous to the basic grinding rule (see §4.1) in that it is generally possible only in marked contexts, but that there are conventional subcases. For example, Atkins (1990) lists:

²²Many adjectives which could normally apply to the animal but which are not usually seen as affecting the meat do not appear in these constructions (??*We serve happy/beheaded chicken* vs. *We serve the meat of happy/beheaded chickens*; see Nunberg, this volume.) However, we think this is explicable on the basis of general pragmatic principles outlined in §5, below.

characteristic dress -> person who wears it (e.g. *blackshirt, red beret*)
musical instrument -> person who plays it (e.g. *cello, sax*).

(Some dictionaries also list, for example, *spear, bow, gun* meaning people who use these weapons, but these seem somewhat archaic.)

Thus we would treat the interpretation of all such novel examples in much the same way as the conventional cases. Novel extended usages are not rare, at least in some styles of writing; (25) is taken from a newspaper travel article.

- (25) [Chester] serves not just country folk, but farming, suburban and city folk too.
You'll see Armani drifting into the Grosvenor Hotel's exclusive (but exquisite) Arkle
Restaurant and C+A giggling out of its streetfront brasserie next door. (Guardian
Weekly, 13 November 1993)

Here *Armani* and *C+A* are presumably intended to be interpreted along the lines of *people wearing clothes from Armani / C+A* (and could be analysed as a combination of two conventionalised processes, brand name -> object, plus characteristic dress -> person who wears it).²³ Our account predicts that all such novel metonymic sense extensions should be analysable as falling into a range of basic patterns which might themselves be language dependent. These basic rules whether conventionalised or not should interact with other grammatical rules appropriately; for example, grammatically induced type coercion occurs when NPs appear as predicative complements, as in (26) (see e.g. Partee, 1992).

- (26) a Sam considers Bill a fool
b Sam is a fool

In (26) *a fool* is coerced from a generalized quantifier to a property (from $\langle\langle e, t \rangle, t\rangle$ to $\langle e, t \rangle$ in extensional terms). Ham sandwich examples can participate in this coercion quite easily, as in (27) said to a waiter delivering a variety of dishes.

- (27) I am the ham sandwich

This is compatible with our account, given that the sense extension will produce a meaning which can be glossed as 'the x who ordered a ham sandwich' which can in turn be coerced to a property of ordering a ham sandwich by the standard type shifting operator.

4.5 Directionality

Although in the case of derivation there is clear evidence of directionality, this is not the case with conversion. In the cases with which we are most concerned where the process is still clearly productive, novel uses, such as the example of *mole* given earlier, at least demonstrate that a particular directionality is possible. In some cases, the basic sense is evident from the morphology, thus we assume that the fruit/nut sense rather than the bush/tree sense is primary in *gooseberry, strawberry, walnut, chestnut* and so forth. This does not preclude the possibility that the direction might change over time nor that there might be cases analogous to morphological back formation. In other cases, there are closely related rules of derivation or compounding which suggest that there should be the same directionality in the conversion case; for example, compounding with *juice* and *meat* closely mirrors the grinding conversion, whilst *-ful* suffixation mirrors the container/contents nominal metonymy. In addition, the tests for cue-validity of transfer functions which Nunberg (1979) proposes can also be used to distinguish basic from metonymic senses, as he suggests, and there appear to be general constraints on transfer functions which suggest that they extend from the concrete to the abstract and the simple to the complex (e.g. Sweetser, 1990).

Cruse (1986:69) describes a test for distinguishing senses according to whether or not they are fully established (i.e. conventionalised in our terminology). This involves the possibility of simultaneously negating the non-fully-established sense whilst asserting the fully established sense, while the converse is much less acceptable. Thus, his example of *novel* meaning the text or the physical object is given in (28).

²³It seems relatively easy to become accustomed to metonymic usages after a particular pattern when they recur in some corpus as though the process were becoming (locally) conventionalised (ham sandwich examples may have this status in the linguistics literature).

- (28) a I'm not interested in the binding, cover, typeface etc — I'm interested in the novel.
 b ? I'm not interested in the plot, characterisation etc — I'm interested in the novel.

It is reasonable to assume that the perceived directionality of sense extension processes would be from fully conventionalised to less conventionalised senses. The examples in (29) seem to confirm the intuition that the animal sense is primary, in cases of meat grinding, and the fruit sense in the fruit tree examples.

- (29) a I don't want the meat, I want the lamb.
 b ?I don't want the animal, I want the lamb.
 c I don't want trees, I want peaches.
 d ? I don't want fruits, I want peaches.

The behaviour in this test is explicable on the assumption that basic conventional senses are assumed by default, and that the extended senses have to be forced by context. There are some cases where neither Cruse's test nor any of the other criterion mentioned give clear results. Nunberg (1978) discusses at length the difficulty of making such a choice in the case of the instance/type distinction, for example. The directionality of sense extension rules does not affect the representation of the signs involved so these preferences in interpretation must follow from the manner of rule application. In §6, we argue that the semi-productivity of such rules can also be used to predict these preferences.

5 Coordination and co-predication

Given that we have suggested two different methods for dealing with systematic polysemy, it is clearly necessary to establish that we can, in fact, distinguish between constructional polysemy and sense extension. It is not always straightforward to distinguish between cases where the relational approach of encoding the different aspects of one entity will work and the examples where it seems necessary to postulate the construction of a new structure via the lexical rule mechanism. Pustejovsky (1994) suggests that the distinction can be made on the basis of co-predication: that *door* can be treated as having a relational structure encoding both the aperture and physical object usages, because of the acceptability of (30).

- (30) John painted and walked through the door.

However, he argues *newspaper* must be coerced between the physical object and organisation usages because of the unacceptability of (31a), despite the acceptability of examples such as (31b) which might be the result of a coercion process applying phrasally to the NP.²⁴

- (31) a * The newspaper fired its editor and fell off the table.
 b John used to work for the newspaper that you are reading.

This is an area where opinions (and judgements) differ. For example, Cruse (1986:65) treats *door* as having distinct panel and aperture senses on the basis of the semantic abnormality of (32):

- (32) ? We took the door off its hinges and then walked through it.

but assumes that a 'global door' sense is involved in (33) (which was cited by Nunberg (1979) as evidence that *door* is not ambiguous).

- (33) The door was smashed in so often that it had to be bricked up.

Care also has to be taken to use cases where the predicates could be true of the same entity, thus (34) does not demonstrate that *teacher* must be coerced.

- (34) ? The teacher was pregnant and had a beard.

²⁴In this particular case, however, it is by no means obvious that *the newspaper that you are reading* has to refer to a physical copy of a paper (see below).

We can assume that acceptable examples of co-predication are evidence that a single structure is available, and thus that constructional polysemy is at work. However, we will argue that zeugma²⁵ is not in general explicable on the basis of the existence of multiple distinct lexical structures. So we cannot necessarily take negative examples as evidence for distinct senses and thus as supporting an account involving sense extension as opposed to constructional polysemy.

In the cases where we have posited rules of sense extension, it is incumbent on us to account for apparent counter-examples involving co-predication. The clearest such examples are coordinations, where there is no possibility of arguing that the sense extension applies phrasally and where the standard rule of coordination requires type compatibility (e.g. Partee, 1992). Furthermore, coordinations involving sortal mismatches are often zeugmatic, as (36) illustrates.

- (36) a He arrived in a Rolls Royce and a temper
 b Our office typist is fast and bearded

Although we argued that similar co-predications of ground and unground senses seem to be ruled out in §4.1, some appear to be possible. For example, (37a) involves a coordination of predicates which select the animal sense of *chicken*, whilst in (37b) we appear to have one which selects both animal and meat senses.

- (37) a This chicken is corn-fed and healthy
 b Corn-fed and inexpensive chicken is difficult to find

We can account for both these examples on the assumption that the ORIGIN of the qualia structure of the ground sense is available for modification (as mentioned in §4.3). Nunberg (this volume) argues that this treatment is insufficiently restrictive since the property described has to have some applicability to the meat for the predication to be fully acceptable. However we would argue that examples such as (38a) are no different from those in which a contextually unexpected adjective is applied straightforwardly to the noun, for example, (38b):

- (38) a ?? We serve corn-fed and happy chicken
 b ? We serve dense potatoes

(38b) is odd, despite the fact that potato tubers can differ in physical density, since it is not generally realised that this affects eating quality. Thus, on our account, both these examples are problematic simply because the context is not providing/supporting a clear interpretation. Making the context explicit improves the acceptability, since it restricts and guides the possible interpretations. Such effects are, admittedly, more likely to arise with adjectives that modify different qualia on our account, but this would be expected, since properties true of aspects of an entity are less likely to relate to a common property, and thus be part of a coherent discourse.

5.1 Coordination in constructional polysemy

In §3 we described a representation of adjectives which relied on selection of predicates from the qualia structure according to the type of the resolved adjectival structure. Adjectives of the same or differing types can be coordinated, although there seem to be some restrictions on the productivity of this process when the adjectives select different qualia (*?fast and well-dressed typist*). But some examples are more acceptable, such as *fast and intelligent typist* where *intelligent* is assumed to be true of the unmodified variable, and the oddness of the others is perhaps better explained as a pragmatic effect. We will assume that the SUBCAT value of the conjoined phrase is the unification of the values on the adjective daughters and that the semantics

²⁵Zeugma is the traditional term for the variety of anomaly which arises when terms are inappropriately linked (yoked) together, such as in (35):

- (35) He was wearing a scarf, a pair of boots, and a look of considerable embarrassment. (from Cruse, 1986:13)

is simply specified as the conjunction.²⁶ Thus we have:

$$[x][fast(e) \wedge P(e, x) \wedge intelligent(x)]$$

where P is coindexed to the telic predicate of the subcategorised noun. This can be applied to *typist* to give

$$[x][fast(e) \wedge type(e, x) \wedge intelligent(x) \wedge typist(x)]$$

Cases such as *corn-fed and expensive chicken* are similar, on the assumption that *corn-fed* selects the ORIGIN in this instance.

Coordination of the noun raises some more complex issues. The first point to notice is that the treatment of adjectives given above precludes the possibility of selecting one role from one conjunct and a different one from another. This appears to be basically correct for adjectival modification. In the example below, *lap* is event denoting and the normal form of *fast* would be expected to apply, whereas it selects for the telic role of *cars*.

(39) ??Prost only gets enthusiastic about fast cars and laps.

This is odd at least on the reading in which *fast* applies to the conjunction *cars and laps*. However, cases where the adjective selects for the same role appear to be generally acceptable, even if the predicate selected differs. For example:

(40) The company's fast typists and computers have raised productivity by 20%.

In such examples, the conjoined entities should be regarded as being combined to produce a single (complex) entity, in order to get the collective readings.

The conjoined form *typists and computers* can be constructed from the individual representations using, for example, the formalism described by Link (1983) to structure the domain such that complex entities can be described. Thus, the semantics of the conjoined phrase could be written as:

$$[x \oplus y][typist(x) \wedge computer(y)]$$

Given the approach that we have adopted previously, of treating the qualia as quite distinct from the rest of the sign, the most straightforward option for the qualia of the conjunction is to identify it with the disjunction of the qualia of the conjuncts.²⁷ In this case, *fast* would select the predicates from the disjunct, giving:

$$[x \oplus y][fast(e) \wedge (compute \vee type)(e, x \oplus y) \wedge typist(x) \wedge computer(y)]$$

But we may want to be able to deduce from this a distributive reading which associates the correct predicate with the particular type of individual (typists who type fast and computers which compute fast). To do this, we would have to complicate the representation somewhat, so that the disjunction was not simply of atomic predicates, but restricted the arguments with respect to the qualia. Although we do not want to equate the fast event with the variables in the qualia structure, we could restrict the fast event to be a subevent of those specified there. In the case of the disjunctive qualia, this would have the effect of restricting fast typing events to the typists and fast computing events to the computers. We will leave this open, since the precise formulation depends on the semantics adopted for events and there are other options, involving alternative treatments of the relationship of the qualia structure to the rest of the sign.

²⁶We will also assume, for the moment, that the type of the conjoined phrase is underspecified. Technically, this raises a problem analogous to that affecting conjunction in HPSG (Pollard and Sag, in press), since the type could not be fully resolved, although, in this particular case, it is possible to define a more complex type system which avoids this situation.

²⁷The main reason why we have maintained the distinction between qualia structure and the rest of the sign here is to avoid making the representations unnecessarily theory dependent. Within HPSG, for example, there are a variety of ways in which the qualia structure might be incorporated into the semantic representation, which would affect the way in which the qualia structure of the conjunct was derived. Qualia could be regarded as part of the BACKGROUND (that is as presuppositional rather than truth conditional) or even be located on the INDEX (Pollard and Sag, in press). These options would carry different implications as to how the qualia should be combined in conjoined phrases. The only essential point here is that the interpretation of examples like *fast typists and computers* where *fast* distributes over the conjuncts requires that the qualia structure of the conjuncts should still be individually accessible in the phrase.

Our current proposal for the representation of verbs like *enjoy*, *begin* and so on, discussed in §3, involves treating them in a manner analogous to *fast*. Conjunctions such as those in (41) are thus possible in much the same way as the conjunction of *fast and intelligent*.

- (41) a Sam picked up and finished his beer
b Sam ate and enjoyed the caviar
c Sam wrote but later regretted that article

However, unlike modification by *fast*, there are some cases where the complement to *enjoy* is a conjunction, such that one conjunct is object denoting and another event denoting.

- (42) a I enjoy films and mending antique clocks
b We found Sam swimming the channel, which he enjoys more than golf (due to Geoff Nunberg)
c Gordon Parry (Gary Mavericks) has come into the world and enjoys a small car, many women possibly including Julia and embezzling the premiums he collects (Guardian, 16th Jan 1990, Features)

In any approach where the ‘coercion’ is internal to *enjoy*, problems arise in treating such examples. No straightforwardly unification based approach can account for both (41) and (42) by postulating one operation applying either to the verb or its complement. If coercion applied to the noun phrase then the noun would need to have a dual coerced/uncoerced nature in (41), if it were internal to the verb then this would have to be both coercing and non-coercing in (42). This remains true even if the work of specifying the coerced meaning is shared between the components, or if the coercion affects part of the sign rather than the whole of it.

Since the examples of conjunction of unlike types in the complement seem more restricted and marked than the conjunction of the verbs, we prefer our current account (which makes (42) problematic rather than (41)) over the one we gave in Briscoe *et al.* (1990) (where the converse applied). The difficulty seems comparable to the problem of cross-categorial coordination from a syntactic viewpoint for which a number of solutions have been proposed (see e.g. Sag *et al.*, 1985; Shieber, 1992; Cooper, 1991). Conjunction is licensed in examples such as (43) if the syntactic descriptions of each of the conjuncts independently unifies with the subcategorisation requirement of the verb, despite the fact that these descriptions will not unify with each other:

- (43) Tigger became famous and a complete snob

Similar remarks must apply to the syntax of examples such as (42) and the semantic effects parallel the syntactic ones: the conjuncts individually have types which are accepted by *enjoy* and the conjunction is only licensed in contexts where *enjoy* (or a similar predicate) is involved. So a promising direction for future research would be to provide an account where this parallelism is explicit. However, any such account will have to move beyond a strictly unification based formalism, to allow for the multiple distinct coercions involved in examples such as (42c).

5.2 Co-predication tests

There are cases where the co-predication test gives less clear indications as to whether constructional polysemy or sense extension are involved. Take the example of *book*: it seems clear enough that it has two senses (or usages) — as a physical entity which represents some text and as the abstract text itself. But the distinction between these is not really straightforward. Consider the set of examples in (44)

- (44) That book is full of metaphorical language.
That book is full of long sentences.
That book is full of spelling mistakes.
That book is full of typographic errors.
That book has an unreadable font.
That book has lots of smudged type.
That book is covered with coffee.

There seems to be a cline here from properties which are clearly true of the content, through those which may be true only of a particular edition or printing through to those which are true only of a copy (cf. Cruse 1986:71). Co-predication of the first and last properties seems odd, as in (45).

- (45) ? That book is full of metaphorical language and is covered with coffee, so it's very hard to read.

But co-predication of adjacent pairs seems natural in all cases, for example (46)

- (46) That book is full of typographic errors and has an unreadable font.

If we treat these senses as cases of constructional polysemy, co-predication is predicted. Thus *book* can have a formal role and a content role in its qualia structure. On this basis, there is no necessary conflict between properties such as *is full of long sentences* and *has coffee spilt on it*. This treatment will not, therefore, account for the apparent oddity of some co-predications. However, although it is standardly assumed that cases of zeugma provide evidence for lexical ambiguity, it is not clear that this is justifiable. Although we must assume, within a unification based account, that acceptable co-predications imply the existence of a single structure, it does not follow that the converse is true. As we suggested above, oddness of co-predication can be simply due to incompatibility of the predicates. Furthermore, there is clear evidence that some sort of pragmatic principle of cohesion must be postulated to account for the unacceptability of some readings where lexical ambiguity cannot be involved. For example, (47) has readings where the gardener bought either fruit or trees, but does not have the crossed interpretations where apple tree and pear fruits were purchased or vice versa.

- (47) The gardener bought three apples and two pears.

Coherence also means that repeated uses of the same homonymous form will tend to have the same interpretation, as in (48), where the crossed interpretation, although possible, is dispreferred (see e.g. van Deemter, 1990).

- (48) John gave four files to Mary and three files to Sue.

Assuming that some such principle is involved, it would also account for the oddness of cases such as *fast and bearded typist*, *tasty and skinny chicken* where we are predicating properties of distinct aspects of the entity, without there being any apparent connection between these aspects. The acceptable examples, such as *fast and intelligent typist*, *tasty and corn-fed chicken*, are those where the distinct aspects are nevertheless related — good typists might be expected to be both fast and intelligent, the food a chicken is given is known to affect the flavour of its meat, and so on (see above). Given this, it is tempting to assume a single structure for *book*.

However, other examples show even more complex polysemy: *newspaper* can also refer to the physical copy or the abstract text (of a particular issue), equivalent sentences to those above can be constructed and the same remarks apply to these as to *book*. But *newspaper* can also refer to an abstract entity other than the text. This is somewhat hard to categorise — it is not necessarily a company, as ownership and editors can change without there being a different newspaper and so on. It seems plausible to suggest that newspapers are regarded as (named) institutions in themselves. Whatever their ontological status, it is clear that in some sentences there is a notion of a 'newspaper-as-institution', but it is not clear that we can make a sharp distinction between this, the content of the newspaper over a number of issues, and the abstract text reading (49).

- (49) That newspaper is owned by a trust.
 That newspaper is left of centre.
 That newspaper supported the Labour Party at the last election.
 That newspaper carries long articles about the internal struggles of the Labour Party.
 That newspaper has obscure editorials.
 That newspaper is full of metaphorical language.
 That newspaper is full of long sentences.
 That newspaper is full of spelling mistakes.
 That newspaper is full of typographic errors.
 That newspaper has an unreadable font.
 That newspaper has lots of smudged type.
 That newspaper is covered with coffee.

Now again, the properties seem compatible with their neighbours, but co-predication of the the first and last is odd, as in (50).

- (50) * That newspaper is owned by a trust and is covered with coffee.

But in some cases co-predication of the copy sense and the organisation sense does seem possible, as in (51) (suggested to us by Geoff Nunberg):

- (51) The newspaper has been attacked by the opposition and publicly burned by demonstrators.

Despite this, assuming that a single structure can cover all the senses of *newspaper* is highly problematic. Constructing a qualia structure to cover all the senses of *newspaper* in such a way that different predicates can apply appropriately is difficult, since it seems that the copy and the organisation sense (at least) should have their own distinct qualia. It is also not clear that one sense can be regarded as primary. Perhaps the most important point is that we can quantify *newspaper* in either the copy or the organisation sense and vagueness of interpretation with respect to the quantification is not possible in such contexts. Thus (52) has the interpretation that three newspapers-as-organisations have been attacked, and some arbitrary number of copies pertaining to each have been burned.

- (52) Three newspapers have been attacked by the opposition and publicly burned by demonstrators.

However, there is no reason within our account why both ambiguity/sense extension and vagueness/constructional polysemy should not be involved, and this would account for the data. Thus for *newspaper*, we assume two structures, one corresponding primarily to the copy and one to the institution. Both of these may be involved in constructional polysemy — the text and parent organisation of the newspaper copy is accessible via its qualia, and conversely the copies are accessible from the structure representing the parent organisation. Note that no intermediate primary structure corresponding to one edition of a newspaper seems to be justified — *three newspapers* cannot mean three editions of the same paper, considered as abstract texts, for example. Thus, in this case, the abstract contents of the physical object can only be accessed indirectly.

Thus the account we have developed here is able to capture facts of co-predication in coordinate structures with constructional polysemy and sense extension insofar as the latter is acceptable. In addition, our account makes further predictions regarding the grammaticality of non-constituent coordination in cases of constructional polysemy. We have not considered the interaction of lexical rules of sense extension with indexical and anaphoric pronouns (see Nunberg, 1993). It is clear that there are many challenges to be faced here, and the consequent complication of the theory of anaphora must be weighed against the advantages gained here in the succinct characterisation of the behaviour of verbs, such as *enjoy*, which subcategorise for multiple complementation within the same or highly related senses, and in the capturing of similarities between sense extension and other lexical processes.

6 The semi-productivity of lexical rules

There are several empirical problems with the account of lexical rules we have developed. Some of these problems are shared with other generative accounts of morphological operations (see e.g. Bauer, 1983 for extensive discussion), others are more specific to our proposal to account for sense extensions in the same fashion. It is well known that morphological processes tend to be semi-productive and are rarely (if ever) exceptionless (e.g. Bolinger, 1975; Aronoff, 1976); for instance, the rule of *-er* nominalisation in English creates deverbal nouns which denote the subject of the underlying predicate – typically an agent, as in *teacher* or *thinker*, sometimes an instrument, as in *(dish)washer* or *(bottle) opener* where the instrumental argument can occur as subject, and occasionally the patient *sticker* or *(best)seller*. However, this rule is not fully productive because items such as *banker* and *stationer* do not have the predicted meaning, whilst a form like *stealer* is blocked by *thief*, though is more acceptable when its meaning is specialised (and made non-synonymous) with a postmodifier – *stealer of fast sports cars / hearts*. Rappaport and Levin (1990) argue that both the agent, instrument and patient versions of *-er* suffixation are rule-governed and the verbs which undergo the latter are at least partly predictable on the basis that they allow middle formation and thus the promotion to subject of the patient argument – *The book sold well*. If we assume that subregularities block regularities and exceptions block all regularities, we can account for this pattern of data without problem. The mechanism required to achieve this looks very similar to that which is required to block *pig* having a meat reading in normal circumstances (Briscoe *et al.*, 1994).

Lexical rules of sense extension, as we have described them, clearly lead to overgeneration. For example, given the sense extension rules for grinding, portioning and animal-metaphor discussed above, (53a) has the interpretations (53b), (53c), (53d) and (53e):

- (53) a John saw some lambs.
b John saw some animals.
c John saw some humans with some lamb-like properties.
d John saw some portions of lamb meat.
e John saw some portions of substance derived from humans with some lamb-like properties.

This problem of rules of sense extension feeding further rules is exacerbated by the lack of morphological marking of the change; that is, the fact that these are rules of conversion rather than derivation. Similar problems arise with uncontroversially ‘morphological’ conversion and derivation; for example, a generative rule-governed approach would have problems explaining why forms such as *unreuntie* are not attested. In the literature on lexical rules, this has led to vacillation between interpretations of lexical rules as ‘redundancy’ statements relating pre-existent entries (e.g. Jackendoff, 1975) and as fully productive generative devices creating new entries from existing ones which match their structural description (e.g. Pollard and Sag, 1987). Neither approach is fully satisfactory since the former fails to capture the semi-productive nature of these rules and the latter leads to overgeneration.

Finally, it is clear that in the case of a sense extension such as grinding, there is distinct variability in the application of the rule to lexical items even within a conventionalised subcase, such as meat grinding; thus, *lamb*, *chicken* and *haddock* are common and established, whilst *mole* and *alligator tail* are not. It is also clear that language users are sensitive to such frequency-based judgements concerning the relative novelty of usages. The same issue arises with derivational morphology in that many forms which are predicted by productive derivational rules are not attested, for example, *hammerer* and *nailer* can be formed by applying *er* nominalisation to the ‘incorporated’ verbs *hammer* and *nail*, respectively. However, English speakers are liable to react to these forms in much the same way they would react to *mole* in the meat sense: with a degree of resistance, but without serious difficulty in interpretation. Bauer (1983:71f), in supporting the view that lexical rules should be treated as fully productive generative rules analogous to those employed in syntactic description, argues that it is this greater ‘item-familiarity’ of lexical items which allows judgements of relative novelty / conventionality to be built up. He points out that there are simply too many combinatoric possibilities at the sentential level for the frequency of particular combinations to be assessed with any confidence by a language user. However, in the case of words and, we might add, idioms the range of possibilities though large is not so great that judgements of novelty based on frequency

lex-count-noun
ORTH = rabbit
CAT = noun-cat
SEM = obj-noun-formula
PROB = 0.4
LRS = grinding(0.05), meat-grdg(0.3), meat-grdg+portioning(0.15), fur/skin-grdg(0.1)

Figure 18: Lexeme for *rabbit*

of use cannot be acquired. Bauer argues, therefore, that accounting for semi-productivity is an issue of performance, not competence.

The frequency with which a given word form is associated with a particular sense (or lexical entry) is often highly skewed; Church (1988) points out that a model of part-of-speech assignment in context will be 90% accurate (for English) if it simply chooses the lexically most frequent part-of-speech for a given word. The incidence of senses of words may well turn out to be similarly skewed. In the absence of other factors, it seems very likely that language users utilise frequency information to resolve indeterminacies in both generation and interpretation. Such a strategy is compatible with and may well underlie the Gricean Maxim of Manner, in that ambiguities in language will be more easily interpretable if there is a tacit agreement not to utilise abnormal or rare means of conveying particular messages. We can model this aspect of language use as a conditional probability that a word form will be used in a specific sense; that is, is associated with a specific entry ($Pr(\text{lexical-entry} | \text{word-form})$). We assume that such probabilities are acquired for both basic and derived senses (lexical entries) independently of the lexical rules used to create derived senses. Thus we make no claim that a derived sense will necessarily be less frequent than a basic one; in the case of a word such as *turkey* in English our intuition is that the ground or animal-metaphor senses are more frequent than the basic sense. It might seem that this assumption commits us to a ‘fully entry’ theory of the lexicon (e.g. Aronoff, 1976) in which all possible words are present; that is, the consequences of lexical rules are precomputed. In the limit, the full entry theory cannot be correct because of the presence of recursive derivational rules such as *re-*, *anti-* or *great-* prefixation in words such as *rereprogram*, *anti-anti-missile* or *great-great-grandfather*, and in our theory of ‘cyclic’ rules of sense extension such as portioning and grinding. Instead we adopt an intermediate position in which we claim that basic entries are augmented with a representation of the attested lexical rules which have applied to them and any such derived chains, where both the basic entry and these ‘abbreviated’ derived entries are associated with a probability.²⁸ For example, a word form such as *rabbit* might be associated with a basic entry like that illustrated in Figure 18, in which meat grinding is shown to be (hypothetically) more probable than grinding, meat grinding and portioning, or fur/skin grinding. Following Cruse (1986) we might refer to this as the lexeme for *rabbit*, in the sense that this basic entry encapsulates our knowledge of the (predictable) behaviour of this word-form (though not of its morphological derivatives, such as *rabbit-like*, and so forth). The attribute LRS associated with the lexeme for *rabbit* records which combinations of lexical rules have been attested with what frequency in the experience of the language user.²⁹ If we assume that speakers choose well-attested high-frequency forms to realise particular senses and listeners choose well-attested high-frequency senses when faced with ambiguity, then much of the ‘semi-productivity’ of lexical rules can be treated as a side-effect of performance. For instance, we would predict that in the ‘null’ or a neutral context (54a) will be interpreted as rabbit meat, and (54b) will be interpreted as animals.

- (54) a John prefers rabbit
 b John wants three rabbits
 c The diners ordered three rabbits

On the other hand, less frequent but attested senses should be chosen when other contextual factors so dictate, as in (54c). In order to specify precisely how this interpretation is preferred, and to formalise the notion of neutral context within this framework, we would need to develop either a thorough-going account of the interaction of lexical probabilities with probabilities associated with specific sentential interpretations, or

²⁸Modulo the probabilistic interpretation, this manner of encoding the (non-)application of a lexical rule has been deployed in many theories; e.g. Flickinger and Nerbonne (1992) and Sanfilippo (1993) in recent accounts of verbal diathesis alternations.

²⁹It is plausible to imagine that language users are able to memorise some estimate of the relative frequency with which a word form and sense occur, though it is unlikely that this process is accurate enough to derive probabilities. Nevertheless, probability theory offers a precise and well-understood theory within which such intuitions can be formalised.

an account of how probabilities reflecting frequency of usage interact with pragmatic principles establishing discourse coherence (or both). This would take us well beyond the scope of this paper, but see e.g. Wu (1990), Lascarides *et al.*, (forthcoming).

In addition to such lexical probabilities, we also think that probability may play a role in the application of lexical rules in novel usage. Under the current proposal, lexical rules will have something akin to the status of ‘redundancy’ rules in that they can be used to create appropriate lexical entries on demand for attested senses of a word form; that is, those which have a non-zero probability in the associated lexeme entry. However, in the situation where an interpretation for a novel usage is called for, an assessment of the relative probability of extant lexical rules would provide a means for adopting the most likely ‘analogous’ interpretation. For instance, interpreting examples such as (55), the listener who had not experienced examples of any variant of grinding with these nouns might choose the rule with the highest probability given the semantic type of the noun.

- (55) a John prefers alligator tail / mole
b John prefers chinchilla
c John prefers pig

The probability of a lexical rule might be derived by comparing the number of lexemes to which the rule could apply (i.e. that it unifies with) where that sense is unattested, to those for which it is attested. Since grinding can apply to any count noun but will be attested for very few, whilst meat grinding can only apply to animal denoting nouns and will be attested for a higher proportion, this predicts that (55a) will be interpreted as cases of meat grinding even in a neutral context. Thus, we can account for productive or ‘analogical’ use of a lexical rule to interpret a novel usage.³⁰ Assuming that the rule of fur/skin grinding is restricted to words denoting animals with fur or ‘good’ skin we may be able to construct a similar account for the preferred interpretation of (55b). However, the notion of semantic type may need to be more fine-grained than is plausible or desirable in a lexicon if we are to account for all such preferences in this manner, since (55b) shows a preference for fur/skin grinding probably as a result of the salience of fur in distinguishing chinchillas from other types of rodent, rabbit or cat. Nevertheless, however this is achieved, it is ultimately a fact about the word and associated sense(s) rather than a fact about animals, since it is irrelevant whether, in reality, more chinchilla animals are worn than eaten. The case of (55c) is different though, since this approach would predict a meat reading on the basis of the greater probability of meat grinding than grinding. However, the preferred interpretation is probably the less specific ‘pig-stuff’ in a neutral context, because of the blocking of this sense by *pork*.

Thus the generation and interpretation of normally blocked forms (unblocking) seems to require a different type of explanation. Briscoe *et al.* (1994) proposed to account for cases of preemption or blocking by introducing a defeasible notion of lexical rule and allowing the output of such rules to be defeasibly overridden in the case where there was preemption by synonymy or by phonological form. The (pragmatic) principle of blocking introduced case specific defeasible blocking statements that could be themselves overridden in pragmatically marked contexts to account for the occasional usages of, for example, *pig* to mean meat with additional affect, and so forth. In this manner, the approach captured Bauer’s (1983:87) insight that blocking is a bar to the institutionalisation (in our terms conventionalisation) of a meaning rather than an outright ban on its use. In this paper, we have presented a rather different formalisation of lexical rules in which the output of the rule itself is not defeasible. From our current perspective, preemption by synonymy can be explained simply by assuming that speakers will use higher frequency forms to convey a given meaning. Thus an extended meaning will not become conventionalised if a common synonym exists. This does not, however, explain the exceptions where blocked forms do occur (except those where the speaker or hearer are unaware of the synonym) nor the effects of their use. The biggest challenge to our current proposal will be to develop an account of the interaction of frequency-based judgements represented as probabilities with default constraints, such as those which allow unblocking. From the perspective of natural language processing a viable alternative might be to model all such pragmatic phenomena probabilistically, perhaps deriving data

³⁰Note that this account has little to say about the conditions under which novel uses will be created, so we will need a further pragmatic theory of the factors licensing novel usage and of the possibility of such usage becoming conventionalised (see e.g. Bauer, 1983). It might be possible to account for the acquisition of lexical rules in terms of a *post hoc* process of generalisation between ‘basic’ and ‘derived’ entries at some point when the productivity of the putative rule reached some probabilistic threshold.

on the frequency of predicted senses from large corpora (e.g. Pustejovsky *et al.*, 1993). However, if we wish to limit the role of probabilities to modelling the frequency-based aspects of semi-productivity and develop theoretical accounts of blocking and unblocking and, say, the interaction of frequency-based judgements with contextual factors favouring a low probability sense, then it will be necessary to utilise a non-monotonic logic in which it is possible to reason about probabilities (see e.g. Pearl, 1988).

7 Conclusion

We have drawn a distinction between some cases of sense modulation and change which we have termed constructional polysemy and sense extension, respectively. This distinction is based on behaviour under co-predication and the traditional distinction between vagueness and ambiguity. We also pointed out in §5.2 that in the absence of clear tests, some cases remain difficult to classify with respect to this distinction.

Both constructional polysemy and sense extension are productive processes which require ‘generative’ lexical mechanisms, in the sense of Pustejovsky (1991). We have proposed to account for some cases of constructional polysemy utilising the notion of nominal qualia structure and predicate coercion. We have formalised this account in a constraint based approach to linguistic description which has been implemented – the LRL/LKB (Copestake, 1992, 1993b). We have argued that this approach, unlike those of Briscoe *et al.* (1990) and Pustejovsky (1993), is capable of capturing many facts of ‘co-predication’. However, our account requires extension in order to deal with the cases of non-constituent coordination discussed in §5.1, in line with other constraint based approaches to coordination (e.g. Shieber, 1992). Furthermore, it needs to be supplemented with a pragmatic account of cohesive co-predication along the lines of Nunberg (this volume), as discussed in §5.

We have argued that sense extensions are semi-productive related sense changes: we cannot simply list all the extended senses in the lexicon, since new ‘analogous’ cases which will not be listed occur. In addition, there are cross-linguistic exceptions and differences of encoding, conventionalised subcases and so forth, which all suggest a sign based, lexical rule account. Nevertheless, sense extensions like other lexical rules of conversion and derivation can be blocked and are applied conservatively. We outlined in §6 an account of the semi-productivity of lexical rules in terms of a probabilistic performance account of their deployment in language production and interpretation. We have also suggested that this account should be integrated with an independent account of blocking or preemption (Briscoe *et al.*, 1994), but this integration remains to be undertaken.

The LRL/LKB framework has also been used to represent cross-linguistic lexical translation (non-)equivalence (Copestake and Sanfilippo, 1993), verbal diathesis alternations (Sanfilippo, 1993) and as a target representational framework for the semi-automatic acquisition of lexical entries from machine-readable dictionaries (see papers in Briscoe *et al.*, 1993 and references therein). In future work, we intend to extend the framework to deal more adequately with default aspects of lexical behaviour and with the integration of lexical and pragmatic phenomena.

References

- Alshawi, H. (ed.) (1992) *The Core Language Engine*, MIT Press, Cambridge, Mass.
- Apresjan, Ju D. (1973) *Regular Polysemy*, Mouton, The Hague, The Netherlands.
- Aronoff, M. (1976) *Word Formation in Generative Grammar*, Linguistic Inquiry Monograph 1. MIT Press, Cambridge, Mass.
- Atkins, B.T. (1990) ‘Lexical Rules: a starter pack’, Ms. Oxford University Press.
- Atkins, B.T. and B. Levin (1992) ‘Admitting Impediments’ in U. Zernik (ed.), *Lexical Acquisition: Using On-Line Resources to Build a Lexicon*, Lawrence Erlbaum, New Jersey.
- Baker, M.C. (1988) *Incorporation: a theory of grammatical function changing*, University of Chicago Press, Chicago.
- Bauer, L. (1983) *English word-formation*, Cambridge University Press, Cambridge, England.
- Bierwisch, M. (1982) ‘Formal and lexical semantics’, *Linguistische Berichte*, **80**, 3–17.
- Bolinger, D. L. (1975) *Aspects of Language*, Harcourt, Brace and Jovanovich: New York.
- Briscoe, E. J., A. Copestake and B. Boguraev (1990) ‘Enjoy the paper: lexical semantics via lexicology’, *Proceedings of the 13th International Conference on Computational Linguistics (COLING-90)*, Helsinki, pp. 42–47.

- Briscoe, E. J., and A. Copestake (1991) 'Sense extensions as lexical rules', *Proceedings of the IJCAI Workshop on Computational Approaches to Non-Literal Language*, Sydney, Australia, pp. 12–20.
- Briscoe, E. J., A. Copestake and V. de Paiva (eds) (1993) *Inheritance, defaults and the lexicon*, Cambridge University Press, Cambridge, England.
- Briscoe, E.J., A. Copestake and A. Lascarides (1994, in press) 'Blocking' in P. St. Dizier and E. Viegas (ed.), *Computational Lexical Semantics*, Cambridge University Press.
- Carpenter, R. (1992) *The logic of typed feature structures*, Cambridge University Press, Cambridge, England.
- Church, K. (1988) 'A stochastic parts program and noun phrase parser for unrestricted text', *Proceedings of the Second Conference on Applied Natural Language Processing (ANLP-88)*, Austin, Texas, pp. 136–143.
- Clark, E. V. and H. H. Clark (1979) 'When nouns surface as verbs', *Language*, **55**, 767–811.
- Cooper, R.P. (1991) 'Coordination in Unification-Based Grammars', *Proceedings of the 5th Conference of the European Chapter of the Association for Computational Linguistics (EACL-91)*, Berlin, pp. 167–172.
- Copestake, A. (1992) 'The representation of lexical semantic information', Doctoral dissertation, University of Sussex, Cognitive Science Research Paper CSRP 280.
- Copestake, A. (1993a) 'Defaults in Lexical Representation' in Briscoe, E.J., A. Copestake and V. de Paiva (ed.), *Inheritance, Defaults and the Lexicon*, Cambridge University Press, pp. 223–245.
- Copestake, A. (1993b) 'The Compleat LKB', ACQUILEX-II Deliverable, 3.1.
- Copestake, A. and E. J. Briscoe (1992) 'Lexical operations in a unification based framework' in J. Pustejovsky and S. Bergler (ed.), *Lexical Semantics and Knowledge Representation. Proceedings of the first SIGLEX Workshop, Berkeley, CA*, Springer-Verlag, Berlin, pp. 101–119.
- Copestake, A. and A. Sanfilippo (1993) 'Multilingual lexical representation', *Proceedings of the AAAI Spring Symposium on Building Lexicons for Machine Translation*, Stanford, CA.
- Cruse, D. A. (1986) *Lexical semantics*, Cambridge University Press, Cambridge, England.
- van Deemter, K. (1990) 'The ambiguous logic of Ambiguity', *Proceedings of the First CLIN Meeting*, Utrecht, The Netherlands, pp. 17–32.
- Fauconnier, G. (1985) *Mental spaces: aspects of meaning construction in natural language*, MIT Press, Cambridge, Mass.
- Flickinger, D. and J. Nerbonne (1992) 'Inheritance and complementation: A case study of *easy* adjectives and related nouns', *Computational Linguistics*, **18.3**, 269–310.
- Godard, D. and J. Jayez (1993) 'Towards a proper treatment of coercion phenomena', *Proceedings of the Sixth Conference of the European Chapter of the Association for Computational Linguistics (EACL-93)*, Utrecht, The Netherlands, pp. 168–177.
- Hale, K. and S. J. Keyser (1993) 'On argument structure and the lexical expression of syntactic relations' in Hale, K. and S. J. Keyser (ed.), *The View from Building 20: Essays in Honor of Sylvain Bromberger*, MIT Press, Cambridge, Mass, pp. 53–110.
- Hobbs, J.R., M. Stickel, D. Appelt and P. Martin (1990) 'Interpretation as Abduction', Technical Note No. 499, Artificial Intelligence Centre, SRI International, Menlo Park, CA.
- Jackendoff, R. (1975) 'Morphological and semantic regularities in the lexicon', *Language*, **51(3)**, 639–71.
- Krieger, H-U. and J. Nerbonne (1993) 'Feature-based inheritance networks for computational lexicons' in E. J. Briscoe, A. Copestake and V. de Paiva (ed.), *Inheritance, defaults and the lexicon*, Cambridge University Press, Cambridge, England, pp. 90–137.
- Krifka, M. (1987) 'Nominal reference and temporal constitution: towards a semantics of quantity', *Proceedings of the 6th Amsterdam Colloquium*, University of Amsterdam, pp. 153–173.
- Lakoff, G. (1987) *Women, fire, and dangerous things: what categories reveal about the mind*, University of Chicago Press, Chicago.
- Lakoff, G. and M. Johnson (1980) *Metaphors we live by*, University of Chicago Press, Chicago.
- Lascarides, A., E.J. Briscoe, N. Asher and A. Copestake (forthcoming) 'Persistent associative default unification', ACQUILEX Working Paper.
- Levin, B. (1993) *Towards a lexical organization of English verbs*, Chicago University Press, Chicago.
- Link, G. (1983) 'The logical analysis of plurals and mass terms: a lattice-theoretical approach' in Bäuerle, Schwarze and von Stechow (ed.), *Meaning, use and interpretation of language*, de Gruyter, Berlin, pp. 302–323.
- Martin, J. (1990) *A Computational Model of Metaphor Interpretation*, Academic Press, Cambridge, Mass.
- Nunberg, G. D. (1978) 'The pragmatics of reference', Doctoral Dissertation, CUNY Graduate Center, reproduced by the Indiana University Linguistics Club.
- Nunberg, G. D. (1979) 'The non-uniqueness of semantic solutions: polysemy', *Linguistics and Philosophy*, **3**, 145–184.
- Nunberg, G.D. (1993) 'On the meaning and interpretation of indexical expressions', *Linguistics and Philosophy*, **16**, 1–43.
- Nunberg, G.D. and A. Zaenen (1992) 'Systematic polysemy in lexicology and lexicography', *Proceedings of Euralex92*, Tampere, Finland.

- Ojeda, A. (1993) *Linguistic Individuals*, CSLI Lecture notes 31. CSLI and University of Chicago Press.
- Ostler, N. and B.T. Atkins (1992) 'Predictable meaning shift: some linguistic properties of lexical implication rules' in J. Pustejovsky and S. Bergler (ed.), *Lexical Semantics and Knowledge Representation. Proceedings of the first SIGLEX Workshop, Berkeley, CA*, Springer-Verlag, Berlin, pp. 87–100.
- Partee, B. (1992) 'Syntactic categories and semantic type' in M. Rosner and R. Johnson (ed.), *Computational linguistics and formal semantics*, Cambridge University Press, Cambridge, England, pp. 97–126.
- Pearl, J. (1988) *Probabilistic reasoning in intelligent systems*, Morgan Kaufmann.
- Pelletier, F. J., and L. K. Schubert (1989) 'Mass expressions' in D. Gabbay and F. Guenther (ed.), *Handbook of philosophical logic, Vol. IV: Topics in the Philosophy of Language*, Reidel, Dordrecht, pp. 327–407.
- Pelletier, F.J., and L. K. Schubert (1988) 'Problems in the representation of the logical form of generics, plurals, and mass nouns' in LePore (ed.), *New directions in semantics*, Academic Press, London, pp. 385–451.
- Pollard, C. and I. Sag (1987) *An information-based approach to syntax and semantics: Volume 1 fundamentals*, CSLI Lecture Notes 13, Stanford CA.
- Pollard, C. and I. Sag (in press) *Head-driven phrase structure grammar*, Chicago University Press, Chicago.
- Pustejovsky, J. (1991) 'The generative lexicon', *Computational Linguistics*, **17**(4), 409–441.
- Pustejovsky, J. (1993) 'Type coercion and lexical selection' in J. Pustejovsky (ed.), *Semantics and the Lexicon*, Kluwer, Dordrecht, pp. 73–96.
- Pustejovsky, J. (1994, in press) 'Linguistic constraints on type coercion' in P. St. Dizier and E. Viegas (ed.), *Computational Lexical Semantics*, Cambridge University Press.
- Pustejovsky, J. and B. Boguraev (1993) 'Lexical knowledge representation and natural language processing', *Artificial Intelligence*, **63**, 193–223.
- Rappaport, M. and B. Levin (1990) '-er Nominals: Implications for the theory of argument structure' in E. Wehrli and T. Stowell (ed.), *Syntax and the Lexicon*, Syntax and Semantics 26, Academic Press, New York.
- Riehemann, S. (1993) 'Word formation in lexical type hierarchies', M.Phil thesis, University of Tübingen, Germany.
- Sag, I., G. Gazdar, T. Wasow and S. Weisler (1985) 'Coordination and how to distinguish categories', *Natural Language and Linguistic Theory*, **3**, 117–171.
- Sanfilippo, A. (1993) 'LKB encoding of lexical knowledge from machine-readable dictionaries' in E. J. Briscoe, A. Copestake and V. de Paiva (ed.), *Inheritance, defaults and the lexicon*, Cambridge University Press, Cambridge, England, pp. 190–222.
- Shieber, S. M. (1992) *Constraint-based grammar formalisms*, MIT Press, Cambridge, Mass..
- Soler, C. and M.A. Marti (1993) 'Dealing with lexical mismatches', ACQUILEX working paper no. 2.4.
- Sweetser, E. (1990) *From etymology to pragmatics*, Cambridge University Press, Cambridge, England.
- Wu, D. (1990) 'Probabilistic unification-based integration of syntactic and semantic preferences for nominal compounds', *Proceedings of the 13th International Conference on Computational Linguistics (Coling90)*, Helsinki, pp. 413–418.
- Young, M. and W. Rounds (1993) 'A logical semantics for non-monotonic sorts', *Proceedings of the 31st Conference of the Association for Computational Linguistics (ACL-93)*, Columbus, Ohio, pp. 209–215.