An approach to building the hierarchical element of a lexical knowledge base from a machine readable dictionary[1]

Ann Copestake
Computer Laboratory, University of Cambridge
June 1990

This abstract describes an approach to extracting taxonomies from machine readable dictionaries and using them to structure a lexical knowledge base which incorporates default inheritance. Taxonomy construction is based on an intuitive notion of the organisation of the substantial quantities of data in machine readable dictionaries which were developed for quite independent purposes. Our intention is to investigate how this affects, and is affected by, the formal semantics of the knowledge representation for the lexical knowledge base which we are attempting to create, especially with respect to inheritance.

1 Introduction

The work described is part of the ACQUILEX project 'The Acquisition of Lexical Knowledge for Natural Language Processing Systems' the goal of which is to demonstrate the feasibility of using the machine-readable versions of conventional dictionaries (MRDs) as a source of lexical information for natural language processing (NLP) systems. We are investigating the extraction of lexical semantic information from the definition fields of lexical entries. In order to do this we have both to define the type of lexical semantic information that NLP systems require and also to determine what information we can acquire (semi-)automatically. The information extracted has to be represented in a lexical knowledge base in a form which can be utilised by an NLP system. Our approach to lexical semantics will determine what information we attempt to extract from dictionary entries, but the content of the dictionaries themselves will in turn affect the theoretical approach, because we view dictionaries as linguistically useful corpora with determinable organisational principles which should be reflected in the knowledge base. However we are really only attempting to represent information that may be derived from the definitions, not the definitions themselves.

The current work is a preliminary attempt to extract information from a MRD and to partially define an appropriate knowledge representation. We have initially attempted to extract taxonomies for three reasons. Firstly taxonomic information appears to be a very basic form of structure inherent in dictionaries. Secondly there is a need for inheritance of information in the approach to lexical semantics which we are currently pursuing. The inheritance hierarchies required seem to have a structure very similar to taxonomies. Thirdly extraction of taxonomic information is known to be feasible; there have been several reports on extracting taxonomies from MRDs (eg Chodorow, Byrd and Heidorn(1985), Calzolari (1988)) and heuristics for establishing the superordinate *genus term* in conventional dictionary definitions are well known. A robust, pattern matching, parser for LDOCE (Longmans Dictionary of Contemporary English) definitions had already been produced (Alshawi 1989). However most previous work on extracting taxonomies has either required lengthy manual preprocessing to disambiguate word senses

[1]This work was supported by Esprit BRA-3030. Thanks to Ted Briscoe, Nicoletta Calzolari and Piek Vossen for their comments on various versions of this paper.

used in definitions or has relied on a user curtailing the process of taxonomy building without giving detailed decision criteria for the curtailment. We have attempted to make sense-disambiguation semi-automatic and to limit the decisions that must be made by the user.

The extracted taxonomies can be regarded purely as data structures; however a hierarchical structure for a lexical knowledge base, the LKB, can be derived from them. Our proposal for a lexical knowledge representation shares much of the motivation behind DATR (Evans and Gazdar 1990), in that it is compatible with unification-based grammar formalisms and allows default inheritance. However our current representation language is considerably simpler than DATR. The taxonomies are viewed as defining a hierarchy down which inheritance can operate. Nevertheless this gives a richer and more flexible system than the sort hierarchies incorporated into the Esprit ACORD system (Moens et al., 1989) and the SRI (Cambridge) CLE system (Alshawi et al. 1989), for example. Our current proposal can be viewed as an augmentation of a standard graph-based unification formalism (such as PATR-II) with minimal default inheritance (allowing only 'orthogonal' multiple inheritance (Touretsky 1986)). However since to some extent we are representing real world knowledge it also seems desirable to consider what the semantics of the hierarchical relationship in terms of real world entities might be.

In the next section I outline a method for semi-automatically extracting sense-disambiguated taxonomies from LDOCE. LDOCE has two features which can be exploited to facilitate this process; the use of a restricted core vocabulary in the definitions makes parsing easier and restricts the possible non-leaf nodes in the taxonomies, and the semantic information encoded in box codes for particular entries can be exploited for sense-disambiguation as explained in Section 2.2. In Section 3 I introduce a proposal for the organisation of the LKB which describes how it can be structured around an inheritance hierarchy derived from the taxonomies. Two potential applications of the LKB are mentioned; in disambiguation and in the treatment of "logical metonymy" (Pustejovsky 1989). I conclude with some preliminary remarks on the semantics of the hierarchical link.

2 Semi-automatic creation of taxonomies.

2.1 The taxonomy creation program

The notion of taxonomy that has been used in work on MRDs is essentially an informal and intuitive one: a taxonomy is the network which results from connecting headwords with the genus terms in their definitions but the concept of genus term is not formally defined; however for noun definitions, which are all I will consider here, it is in general taken to be the syntactic head of the defining noun phrase. The program described, in effect, implements this definition with the further (incorrect) assumption currently being made that a genus term is a single lexeme (coordinated genus terms are discussed in Section 3.2).

The taxonomy creation program operates in a top-down fashion; starting from a given, disambiguated word sense all the entries in which this is used as the genus term are found, and the program then recurses on the word sense corresponding to each of these entries, terminating on word senses which are never used as genus terms. So starting from **liquid 2 (1)**, **oil 1 (1)** is one of the (large number of) entries found; this in turn is the genus term in entries such as **paraffin 1 (1)**. **Paraffin 1 (1)** is not used as a genus term in any entries, so the process bottoms out at this point, and **paraffin 1 (1)** will be a leaf node in the taxonomy.

The parser retrieves the *class* from definitions of individual word senses; this could be lexically ambiguous with respect to both homonyms and subsenses of the genus term. Most of our experiments have been done on definitions of nouns — for these the class will normally be the syntactic head of the noun phrase, the exceptions being definitions which include phrases such as "a kind of X" where the class assigned will be X rather than "kind". For details of the parser see Alshawi (1989). Finding the entries with a given genus term is thus a two stage process; first all the entries with the appropriate class will be found and then these are filtered by disambiguating the class. For example, in order to find the entries where **liquid 2 (1)** is used as the genus term, the entries which have "liquid" as a class are first retrieved and these are then examined to attempt to determine whether **liquid 2 (1)** is the sense being used, as opposed to one of the adjectival senses under **liquid 1**, or the technical sense **liquid 2 (2)** "either of the consonants r or l". A fragment of the taxonomy derived starting from **liquid 2 (1)** is shown below.

- liquid 2 (1)
  - drink 2 (1 2)
    - absinth 1 (1)
    - anis 1 (1)
    - aperitif 1 (1)
    - applejack 1 (1)
    - arrack 1 (1)
    - barley water 1 (1)
    - beer 1 (1)
      - ale 1 (1)
      - barley wine 1 (1)
      - bitter 3 (1)
      - bock 1 (1)
      - home brew 1 (1)
      - lager 1 (1)
      - light ale 1 (1)
      - mild 2 (1)
      - pale ale 1 (1)
      - porter 3 (1)
      - stingo 1 (1)
      - stout 2 (1)
      - wallop 1 (2)
    - beer 1 (2)
    - beer 1 (3)

Disambiguation is essential in order to produce a coherent taxonomy and to allow termination of the taxonomy building process before it potentially includes all the entries in the dictionary. For example, one of the entries which has **liquid** as a class is **spirit 1 (13)**. If the program recursed on **spirit** and did not do the sense discrimination (or did not do it correctly) entries such as **devil** and **ghost** would be retrieved. These in turn would give rise to further problems of the same kind.

In the case of **liquid**, disambiguation will be quite straightforward; the adjectival homonym is excluded since the taxonomy is restricted to nouns because of the convention that a sense is defined by a phrase of the same part of speech (so the genus terms in definitions of nouns will always be nouns, for example). A set of (weighted) heuristics is used to find the "best match" among the possibilities with the same syntactic category as the potential genus term — in the example **liquid 2 (1)** and **liquid 2 (2)** are thus the only candidates.

The current set of heuristics rely partially on a feature particular to LDOCE — most entries have a *box code* in the MRD which encodes a certain amount of semantic information. In particular one part of the box code contains codes such as L for liquid, H for human, T for abstract. One heuristic is based on whether the value of this code for the potential genus sense is compatible with the code for the entry. For example, the codes for **oil 1 (1)** and **liquid 2 (1)** are both L, but that for **liquid 2 (2)** is T, which is not compatible with L. Another part of the box code specifies whether a sense is slang, informal, technical and so on. Slang and informal senses in particular are very unlikely to be used as genus terms. Another heuristic is simply that lower numbered senses are more likely to be used as genus terms, because the ordering of senses tends to reflect their relative frequency/typicality (although there are many exceptions). A fourth heuristic attempts to determine the degree of overlap between the uninflected forms of the words in the definitions, to provide some measure of semantic "closeness".

Such sense disambiguation is, of course, not completely reliable. So taxonomy derivation is semi-automatic — the user is asked confirm or override the important choices, those concerning the entries which might be non-leaf nodes in the taxonomy. (In our tests so far this has resulted in the user being consulted on about 5% of the entries.) This is enough to prevent spurious paths being followed, furthermore an incorrect decision about a leaf node cannot affect the final taxonomy too greatly. The user also has the option to specify that the parser has assigned the class term incorrectly. Such yes/no decisions are normally fairly quick and easy to make. The only other decision that the user has to make is whether to merge subsenses which might be too difficult to disambiguate. For example **animal 1 (1)** and **animal 1 (2)** might be amalgamated because **animal 1 (2)** is defined as the group of **animal 1 (1)** excluding humans. (The new sense is specified as **animal 1 (1 2)**.) [2]

2.2 Test results

The program has mainly been tested by building two substantial taxonomies, one starting with **animal 1 (1 2)** and one with **substance 1 (1)**. The animal taxonomy includes about 700 words senses, of which only about 30 are non-leaf nodes and has a maximum depth of 6; the substance taxonomy has about 1000 senses with 50 non-leaf nodes and a maximum depth of 5. Taxonomies derived from LDOCE would be expected to have a greater percentage of leaf nodes than those from other dictionaries because of the restricted core vocabulary.

Together these taxonomies cover about 7% of the noun entries in LDOCE yet it took less than two hours of interaction to derive them both. Assuming that these examples are fairly representative it seems that it will be possible to construct taxonomies containing the majority of noun senses relatively quickly. An important side effect of the process of taxonomy building is the disambiguation of the genus terms in the definitions; we expect to be able to exploit this information when parsing entries to extract other information.

No distinction is made in these taxonomies between mass nouns, group nouns and singular count nouns or between senses which seem to correspond to natural kinds and those which do not. Although it is clear that some of this information can be extracted from the entries we currently ignore it in producing the taxonomies. I do not think that there should be separate taxonomies for count, mass and group nouns. For example, in definitions of the form "a group of X" we take X to be the genus term rather than **group**. Although the information that something is a group term, for example, should be represented in the LKB, it seems inappropriate to represent it in the taxonomy. To do this extra link types would have to be invented, but this seems likely to result in the imposition of a structure on the entries which is not really there because it complicates the concept of genus term as accepted by lexicographers. Representation can really only be considered with respect to a particular linguistic treatment. I developed an approach to the semantics of mass terms and plurals in Copestake (1990a) (after Krifka 1987) where singular, group and mass entities have the same model-theoretic status and the distinctions between them are expressed using a classifier-like construction, which seems

[2] Merging of senses has been avoided wherever possible in our experiments so far but will be necessary, so we need to define more stringently when and where it is justified.

to justify the use of a single taxonomy (this is considered in more detail in Copestake (1990b)).

Although there were some erroneous inclusions in the taxonomies due to the sense discrimination heuristics failing, most of these were relatively unimportant, because they were failures to discriminate between senses which were fairly similar. An exception was the erroneous inclusion of 6 senses such as **crew 1 (3)** under **team 1 (1)** in the animal hierarchy. **Team 1 (1)** is defined as "2 or more animals drawing the same vehicle" whereas the group of people sense is **team 1 (2)**. However this sort of error is very easy to spot, because there are a group of incorrect words — manual checking/editing of the complete taxonomy to correct this sort of error would be relatively easy.

The problems caused by Alshawi's parser and grammar failing to identify the class term in the definition correctly were much more serious. Up to 20% of the senses which were assigned class ANIMAL should have been assigned some other class and over 150 word senses were omitted because the class assignment was missing or wrong[3] Many of these cases were in definitions where there was a conjunction eg "a person or animal who ..." which the grammar does not cover. Such definitions raise questions about the structure of the taxonomy which are discussed in section 3.2. In other cases there was structural ambiguity and the (single) result returned was incorrect:

armadillo "a small animal native to the warm parts of the Americas" — was assigned class NATIVE.

We are currently attempting to build a more sophisticated parser for LDOCE definitions, which should solve some of these problems.

There are problems which are inherent in the dictionary entries. There are cycles in the definitions. For example **animal 1 (1)** is defined as "a living creature ..." and **creature 1 (1)** is defined as "an animal of any kind". Because the taxonomy was produced starting from **animal** this does not cause any serious problems; **creature 1 (1)** is included in the hierarchy under **animal 1 (1 2)**, recursion on **creature 1 (1)** happens as normal except that the cycle is detected and noted. A more serious problem is when the cycles result in words being excluded altogether — for example **cow** is defined in terms of **cattle** and **cattle** in terms of **cow** — about 17 word senses were excluded from the animal taxonomy because of this. In general this could only be detected by creating taxonomies which covered the majority of the entries and then determining which class words had not been incorporated.

3 Structuring the lexical knowledge base with taxonomic information

3.1 The knowledge representation

One view of the LKB is as a lexicon with links between partially specified entries. Inheritance of information about particular attributes along these links results in complete entries which could then be regarded as independent entities equivalent to entries in a normal lexicon. A conventional lexicon produced by expanding all the entries at once with the inherited information would probably be unmanageably large — the point is that the

[3] We checked for omissions by retrieving all the entries with the appropriate box codes for animals, but this will not be possible in general.

links are only seen as providing for inheritance. We currently represent the attribute structure for individual entries as labelled DAGs, giving expanded entries similar to those used in PATR-II. I will initially concentrate on this view of the LKB; the definition of inheritance in such a system is analogous to that in an object oriented programming language — it specifies what attributes an LKB entry will have but does not constrain its interpretation. However it is also possible to see at least some of the contents of the LKB as real world knowledge which should be given a semantics in terms of real world entities; some preliminary remarks about the semantics of the links in these terms are made in section 3.3.

The LKB, for noun senses at least, will contain links derived from the taxonomies, that organise the entries as a tree-structured hierarchy. I will refer to this as an IS-A hierarchy, to distinguish it from other possibilities such as a PART-OF hierarchy, but I am using IS-A in a somewhat broader way than usual because mass and group terms will be included in the hierarchy. The earlier assumption of a unique genus sense means that our taxonomies in general will be tree-structured (exceptions are discussed in 3.2). This is a very useful property because it gives a unique ordering for a default mechanism; we have therefore assumed that the IS-A hierarchy can be tree-structured. The problems of allowing defaults in systems which incorporate multiple inheritance are well known (the so-called "Nixon diamond" is an example). Since it greatly simplifies our task we start from the assumption that only orthogonal multiple inheritance will be allowed (see Touretzky 1986). Thus if we do add other relationships such as PART-OF, any information inherited along PART-OF links would not be allowed to conflict with that inherited along IS-A links. Under these conditions inheritance has been formalised in a variety of ways (Touretzky 1986, see also Evans and Gazdar 1990); since the appropriate behaviour seems uncontroversial I will not discuss it further here.

One straightforward use of the LKB is to allow features such as ANIMATE to be represented for use by a parser which incorporated selectional restrictions, for example. Such features could be associated either manually or semi-automatically with the LKB entries corresponding to the nodes towards the top of the taxonomy; the more specific word senses would then inherit them. However the approach to lexical semantics which we are currently pursuing requires the inheritance of a much richer range of attributes. For example the treatment of logical metonymy adopted in Briscoe, Copestake and Boguraev (1990) allows the coercion of sentences such as:

John enjoyed the autobiography.

into the equivalent of:

John enjoyed reading the autobiography.

(in an informationally impoverished context). Assuming that the information that the default predicate associated with **book 1 (1)** is **read 1 (1)** has been acquired in some manner (currently manually, possibly eventually through analysis of dictionary definitions or corpora) this information would then be inherited by the entry for **autobiography**. Entries lower in the taxonomy may override the inherited information; thus although **telephone directory** is under **book 1 (1)** in the hierarchy the reading event should be cancelled by more specific information about the purpose of telephone directories.

We might represent this sort of information in the LKB as follows:

book 1 (1)	[default: purpose : read 1 (1) animate: false]
autobiography	[IS-A: book 1 (1)]
telephone directory	[IS-A: book 1 (1) default: purpose : refer to 1 (2)]

The IS-A links are to be interpreted as allowing inheritance, once this has occurred the expanded entry can be regarded as a normal entry in a lexicon. **autobiography** would thus have the same values as **book 1 (1)** for both attributes. **telephone directory** will inherit the value false for the attribute animate but the value of its purpose attribute overrides that inherited from **book 1 (1)** — the assumption is that failure of ordinary unification will always lead to the information associated with the node lowest in the hierarchy being retained, regardless of any possible relationship between the values.

3.2 Constructing an IS-A hierarchy from taxonomies

In order to build a single hierarchy, the taxonomies derived by the methods described in the previous section would have to be connected at the top. This would have to be done by hand, since dictionary definitions for very general terms such as **substance** and **animal** are not very helpful; to some extent at least the top nodes will be theory-dependent and will not correspond directly to word senses. However the dictionary generated taxonomies can be expected to constrain this top level structure. The cycles that exist in the definitions should be removed in some way. This could be done by making a (possibly arbitrary) decision to break one link or by specifying that the two senses are synonyms (and therefore occupy the same position in the taxonomic ordering and have compatible values for the relevant attributes).

Other researchers have assumed that a tangled hierarchy is needed to represent the taxonomy (see Chodorow, Byrd and Heidorn 1985). In most cases this is because senses have not been disambiguated. The other cases are entries with coordinated genus terms. As mentioned earlier the current parser does not handle these but our intention is to represent them in the taxonomy as structured rather than atomic nodes. Representation in the LKB depends on the conjunction.

For example, coordination with "or" such as in:

hunter 1 (1) a person or animal who hunts ...

is very common in LDOCE. In some definitions it seems as though two word senses are being combined; in order to represent **hunter 1 (1)** in a tree-structured hierarchy the sense could be split:

hunter 1 (1a) a person who hunts ...

hunter 1 (1b) an animal which hunts ...

but this does not really seem justified. Clearly having a tangled hierarchy does not solve the problem; **hunter 1 (1)** cannot have both **person** and **animal** as parents in any inheritance hierarchy because it would then be given the union of their attributes — it would be a chimera of person and animal. Instead **hunter 1 (1)** should inherit the intersection of the attributes of **person** and **animal** — this might best be implemented by specifying an LKB entry **person or animal**[4] since this would apply to over 70 definitions.

In contrast, in order to correctly represent symmetric coordination by "and" in the genus term, multiple inheritance would be required. However there are relatively few examples of genus terms containing "and" in the noun definitions and all those which I am aware of seem to involve a temporal use — for example:

montage the choosing, cutting and combining together of separate photographic material ...

We do not currently have any way of adequately representing this in the LKB because we have only been considering the representation of static information, but simple multiple inheritance from the conjuncts is clearly not what is required.

3.3 Inheritance and the semantics of the IS-A link

Although it is appropriate to describe the semantics of a lexicon in terms of how attributes are inherited, without constraining the interpretation of the entries, this seems inadequate for the LKB. The taxonomic information that we are attempting to represent involves real world knowledge; furthermore, since we are extracting it automatically from MRDs, we ought to attempt to justify its use and the inheritance behaviour. The following is discussed in more detail in Copestake(1990b).

The simplest possibility is that the IS-A link in the LKB should be interpreted as specifying that a subset relationship exists between the extensions of the words senses, for example:

$$\text{autobiography ISA book1(1)} \equiv [\![\text{autobiography}]\!] \subseteq [\![\text{book1(1)}]\!]$$

In good lexicographic practice the genus term should "pinpoint that property of a thing that is viewed by most speakers as being essential to it" (Landau 1984). Thus we might expect that the subset relationship is (usually) entailed by the taxonomic relationship (whatever its actual semantics are). Our results so far are mostly consistent with such an interpretation.

If this interpretation of the IS-A link in the LKB is adopted then inheritance of necessary properties will not be overridable. If it is necessarily the case that animals are animate then everything below animal in the hierarchy must inherit the property animate. Default inheritance is needed to allow inheritance of default properties. A possible way to formalise this involves making use of Delgrande's (1988) default logic. This allows the representation of statements such as:

$$\Box(x)[\text{book(1)1}(x) \supset \neg\text{animate}(x)]$$
$$\Box(x)(\text{autobiography}(x) \supset \text{book(1)1}(x))$$
$$(x)(\text{book(1)1}(x) \Rightarrow \text{purpose}(x, \text{read1(1)}))$$

(Autobiographies are necessarily books, books are necessarily not animate, normally the purpose of books is to be read.) From this it follows that:

[4]In this case **person or animal** actually corresponds to the word sense **animal 1 (1)** but in general we would need a small number of nodes in the IS-A hierarchy that did not correspond to particular word senses.

$$\Box(x)[\text{autobiography}(x) \supset \neg\text{animate}(x)]$$

but more interestingly one of Delgrande's approaches to default reasoning would involve concluding by default:

$$(x)(\text{autobiography}(x) \Rightarrow \text{purpose}(x, \text{read1}(1)))$$

This seems to give a justification for the LKB's proposed default inheritance mechanism which would link in with a mechanism for reasoning about particular individuals by default.

There are definitions which do not straightforwardly fit into this approach. For example:

water 1 (1) the most common liquid ...

ice 1 (1) water which has frozen to a solid ...

We could represent this in the LKB as:

liquid 2 (1)	[state: liquid 1 (1)]
water 1 (1)	[composition: water 1 (1) IS-A: liquid 2 (1)]
ice 1 (1)	[state: solid 1 (1) IS-A: water 1 (1)]

and allow the attribute state to be overridden by more specific information. But this implies that **liquid 2 (1)** is to be interpreted as something that is by default liquid. I don't actually find this unreasonable but we then have to explain why:

Ice is a liquid.

seems false. Giving a real world semantics to the IS-A link makes such problems apparent, and may give us some insight into the nature of the taxonomic relationship; these examples tend to suggest that our implicit assumption of a static taxonomy is problematic, for example.

However we believe that the correct approach is not to speculate about possible solutions or to suggest a more complex approach at this stage but to derive further taxonomies and to attempt to classify the problems before refining our approach. Since we are primarily attempting to specify the semantics of the LKB and not that of the taxonomic relationship, if we can identify the definitions that cause particular problems it may be possible to circumvent them in the LKB by, for example, introducing entries that do not directly correspond to particular word senses, although this is obviously something to avoid if possible. The frequency and nature of the exceptions would determine whether we do this or modify the semantics.

4 Conclusion

The work described is at a very preliminary stage. We intend to extract more taxonomies and to build an LKB which is adequate for representing large amounts of information. Exactly what attributes we attempt to extract and represent will depend on the problems that we tackle as part of the ACQUILEX project.

Besides the issues mentioned here concerning the representation of mass and group terms, the semantics of the hierarchical relationship in the LKB, and of the taxonomic relationship itself there are a variety of other questions connected with the extraction of taxonomic information which will have to be addressed eventually; for example the way in which sense distinctions are and should be made (see Vossen 1990), and the status of natural kinds in the taxonomies. The use of a large knowledge base representing information extracted from MRDs will be a fundamental part of our approach to these issues.

References

Alshawi H(1989) 'Analysing the dictionary definitions' in Boguraev B and Briscoe E J (eds.), *Computational lexicography for natural language processing,* Longman, London, pp.153–169

Alshawi H et al(1989) *Research programme in natural language processing, final report,* SRI Cambridge Computer Science Research Centre, Cambridge

Briscoe E J, Copestake A A and Boguraev B K(1990) *Enjoy the paper: Lexical semantics via lexicology,* Coling 90 (in press)

Calzolari N(1988) 'The dictionary and the thesaurus can be combined' in Evens M W (eds.), *Relational models of the lexicon,* Cambridge University Press, pp.75–96

Chodorow M S, Byrd R J and Heidorn G E(1985) 'Extracting semantic hierarchies from a large on-line dictionary', *Proceedings of the 23rd ACL,* University of Chicago, pp.299–304

Copestake A A(1990a) *Some notes on Mass Terms and Plurals,* Technical Report 190, Computer Laboratory, University of Cambridge

Copestake A A(1990b) *The representation of lexical semantic information,* in preparation

Delgrande J P(1988) 'An approach to default reasoning based on a first-order conditional logic: revised report', *Artificial Intelligence, vol.36,* 63–90

Evans R and Gazdar G (eds)(1990) *The DATR papers,* Cognitive Science Research Paper CSRP 139, School of Cognitive and Computing Sciences, University of Sussex

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

Landau S I(1984) *Dictionaries: the art and craft of lexicography,* Scribner, New York

Moens M et al(1989) 'Expressing generalisations in unification-based grammar formalisms', *Proceedings of the 4th European ACL,* Manchester, pp.174–181

Pustejovsky J(1989) 'Current issues in computational lexical semantics', *Proceedings of the 4th European ACL,* Manchester, pp.xvii–xxv

Touretzky D F(1986) *The mathematics of inheritance systems,* Morgan Kaufmann, Los Altos

Vossen P(1990, forthcoming) 'Polysemy and vagueness of meaning descriptions in the Longman dictionary of contemporary English' in Svartvik J and Wekker H (eds.), *Topics in English Linguistics,* Mouton de Gruyter, Amsterdam