

# An Overt Semantics with a Machine-guided Approach for Robust LKBs

Evelyne Viegas  
New Mexico State University  
Computing Research Laboratory  
Las Cruces, NM 88003  
USA  
viegas@crl.nmsu.edu

## Abstract

In this paper, we report on our experience in building computational semantic lexicons for use in NLP applications. In a machine-guided approach, the computer induces part of the semantic knowledge to be acquired by an acquirer. An overt semantics can help predict the syntactic behavior of words. By overt semantics we mean applying the linking or lexical rules at the semantic level and not on lexical base forms. More specifically, we address the different strategies of acquisition arguing for an application-driven, training-intensive effort. We also report on how to develop lexicons using off the shelf resources, and address multilingual issues. We will try to provide an assessment of the difficulties we encountered and some directions to bypass them.

## 1 Introduction

Our experience in building computational semantic lexicons which are used by Natural Language Processing (NLP) systems comes from Mikrokosmos, a knowledge-based machine translation system,<sup>1</sup> where texts from Spanish and Chinese are translated into English. Mikrokosmos adopts an interlingua-based approach (Nirenburg *et al.*, 1992) and all lexicons can be used for multilingual analysis and generation: each word is mapped to an interlingua structure. The lexicons built for Mikrokosmos are multi-purpose **multilingual** to support translation or multilingual generation tasks, **reusable**, that is, applicable to several NLP tasks, (e.g., generation, analysis, information extraction), and **maintainable**, that is, supporting semi-automatic acquisition and restructuring of the lexicons.

The content of the Lexical Knowledge Base (LKB) is essentially the same irrespective of a particular application. The types of information important for analysis and generation might differ, as suggested by Dale and Mellish (1998). For instance, recording all the senses of a lexeme is more important for analysis than generation, conversely, knowing stylistic information on words such as *highfalutin* or *formal* is

<sup>1</sup>For a description of Mikrokosmos, see <http://crl.nmsu.edu/Research/Projects/mikro/index.html>

important for generation (Hovy, 1988). The content of a multi-purpose LKB is application independent (modulo its indexing: in analysis the LKB is indexed on lexemes whereas for generation the LKB is indexed on concepts). We argue, in section 2 that the acquisition process is application-dependent. Moreover we argue that defining the meaning of a word for NLP systems requires a training-intensive effort. In other words, the fact that we, as humans, understand texts does not entail that we can determine the “computational” meaning of a word. Chomskian trees are linguists’ constructs, not innate structures. A linguist must be trained to be able to build syntactic patterns (e.g., trees). In computational semantics, the same rule applies: one must be trained to build the corresponding semantics (e.g. frames, predicates, ...) for a word. In order to approach the “computational” meaning of a word, training is the most important means we have to date to ensure consistency among acquirers. Other means are to adopt an overt semantics with a machine-guided approach which directs as much as possible the acquirer (Section 3). This machine guided approach could also act behind the ‘back’ of an acquirer “correcting” some inconsistencies in lexical descriptions between acquirers, as will be shown in Section 6. In Section 4, we discuss our use of off the shelf resources, such as WordNet (Miller, 1990), to accelerate the machine-guided acquisition of the English lexicon by taking advantage of the existing database of synsets<sup>2</sup> which provide synonym lists for a lexeme. We also show how a semantic-based approach, can help predict the syntactic behavior of words. Note that the reverse (predicting semantics from syntax) is not true, as some experiments on Levin’s work (1993) have shown (Section 4). In Section 5, we address multilingual issues in lexicon development.

## 2 Application-driven Acquisition

The semantics of an entry is an underspecified Text Meaning Representation (TMR) fragment (e.g., De-

<sup>2</sup>Synsets represent WordNet’s building blocks which are words, synonyms or near-synonyms, that can be used to refer to a given concept (Miller, 1990)

frise and Nirenburg, 1991) This TMR fragment can be a concept from the ontology or some interlingua structures such as attitudes, modalities, aspects, sets and TMR relations (addition, enumeration, comparison) Concepts and interlingua structures can appear together or independently The ontology, to which lexemes are mapped, consists of concepts (named sets of property-value pairs) organized hierarchically along subsumption links, with an average of 14 relational links (such as ISA, SUBCLASS, AGENT, THEME-OF, HEADED-BY, HAS-MEMBER) per concept (Mahesh, 1996) In a multilingual environment, the main practical advantage of connecting the lexicon to an ontology is cost-effectiveness, as only the "language-dependent" properties have to be acquired when adding new natural languages to the system

The mapping between a word and the ontology is the most difficult task of lexicon acquisition, and requires to develop the most cost-effective approach in terms of training and strategies

## 2.1 Importance of Training

The experiment reported below shows that training is essential to determine the "computational" meaning of a word A native speaker of Spanish, who had not taken part in the lexicon training process, was asked to add some senses to entries in the Spanish lexicon This was mainly done for testing the analyzer, as there were only 23 out of 167 words which were ambiguous in one text we were analyzing But we also discovered this was a very useful exercise for testing the quality of a semantic lexicon

The list of added senses was reviewed by two computational linguists, one in charge of supervising the training and the other with proficiency in our framework who had seen entries as they were used by the analyzer but had not taken part to the training process either The untrained acquirer, hereafter UNACQ, added a total of 111 to 55 open class words or so Among these 55 words where ambiguity had been added, 33 were already ambiguous in the Spanish lexicon

After a closer look at the Spanish lexicon, and at the senses retrieved by the semantic analyzer, and after doing an on-line corpora search, the computational linguists accepted less than 20 new senses among the 111 suggested This "overgeneration" of senses by UNACQ had different origins 1) the analyzer did not present all the senses from the Spanish lexicon to UNACQ, it only presented the ones that were accepted after syntactic binding, 2) the senses added by UNACQ were "equivalent" to the senses already in the Spanish lexicon, but not recognized by UNACQ, as they were acquired as "unspecified" in the Spanish lexicon, 3) UNACQ hard-coded non-literal meanings of the words, 4) the addition of senses was MRD-driven UNACQ acquired

the list of meanings provided by the Spanish-English Larousse and Collins, adopting an enumeration approach Such a task is not superficial, it ensures that the quality of the core lexicon is good enough so that it can serve as a basis for lexicon expansion techniques, some of which we develop below (see Viegas (1999) for the choices an acquirer faces when working out the semantic mapping of a word)

## 2.2 Strategies

There are mainly two approaches to word sense assignment corpus-driven and mental-driven The former is better adapted to building lexicons used in analysis, whereas the latter better suits lexicons to be used in generation We refer to Kilgarriff (1997) for the corpus-driven approach, and discuss in this paper the mental-driven approach A mental-driven or thesaurus-driven approach consists in grouping together lexemes which share the same meaning In order to ensure consistency among acquirers' mappings we have divided the process of acquiring a computational semantic lexicon into two phases pre-acquisition and acquisition There is still time to revise a pre-acquired mapping at acquisition time, if needed

## 2.3 The Pre-Acquisition Phase

For a generation lexicon, the method of preparing the pre-acquisition files can be as follows 1) extract all concepts from the ontology, 2) lexicalize them using on-line thesauri, dictionaries and native speakers' intuitions, 3) order pre-acquisition files according to the semantic Mapping-Tag (see below)

A pre-acquisition record includes 7 fields Semantics, Mapping-Tag, Lexeme, POS, Translations, Frequency, and Polysemy-Count

The Semantics field includes only the ontological head concept, in which the word sense should be anchored (no selectional restrictions or other properties are specified at this stage) The Mapping-Tag field (see below) describes the type of connection between the word sense and its conceptual meaning some word senses are directly mapped ("dim" map) to a single concept in the ontology, whereas the meaning of some other word senses is described through the combination of concepts linked via properties (relations or attributes) We defined seven tags which flag the entry for a specific task For instance, "devb" (deverbal) is used primarily for nouns and adjectives when their meaning is a composition of a filler and an event (e.g. *bombing*, *readable*), "asp" (aspectual) is used for true aspectuals (e.g. *begin*) and also with actions expressing aspectuality (e.g. *stare*, duration prolonged) The Translations field includes an English translation (for languages other than English) Frequency, POS and polysemy count are extracted automatically, using on-line large corpora for frequency, and WordNet for

the part of speech and the polysemy count. Bilingual dictionaries, filtered by a native speaker of the foreign language, are used for the translations into English (for the acquisition of languages other than English)

In order to increase speed at acquisition time, each acquirer works on one type of Mapping-Tag at a time. For instance, some acquirers work on type OBJECT. Type OBJECT can only be lexicalized into nouns, e.g. DEVICE → *device instrument tool appliance*. Others work on the type EVENT. EVENTS can be lexicalized into nouns, e.g. EXPLODE → *bombing, bombardment*, or into verbs *bomb, bombard, drop bombs on, throw bombs at*. In order to increase consistency, acquirers go through specially designed training sessions.

### 3 Overt Semantics to Predict Syntax: A Machine-guided Approach

Mappings between semantic roles and syntactic complements are defined via a mapping (a rule). These mappings can be defined for large sub-classes of lexical entries. For example, the rule Att-Pred-Adj creates an entry which accepts in the semantic feature a concept from the subtree of ATTRIBUTE or an ATTITUDE and accepts attributive (e.g. *safe car*) and predicative uses (e.g. *the car is safe*). In the case of an adjective mapped to a RELATION (e.g. MENTAL-OBJECT-RELATION) the preferred rule would be Att-Adj generating an attributive reading (e.g. *dental practice*), and not (*\*the practice is dental*)

By selecting the appropriate mapping for classes of entries, it is possible to hide the mapping from the acquirer since these mappings are defined in a lexical class, not in an instance. As defined by an acquirer, an entry looks as follows

```
[key "safe",
 syn Att-Pred-Adj,
 sem [name Safety-Attribute,
      range Safe]],
```

During compilation of the dictionary, the Att-Pred-Adj label is replaced with its definition and makes explicit the co-reference between the subcategorization and the semantics.

So far we have developed for the English lexicon about twenty syntactic patterns which apply to a large number of semantic frames. In the case of adjectives, we have 3 rules, one for attributive adjectives, another one for predicative adjectives, and a third one for attributive adjective used predicatively.

In the case of nouns, we have developed four patterns as illustrated below

Subcat	pattern	Example
N		weapon
NObl10pt		father (of two)
NObl1-Ob120pt		bombing of Iraq (by the US)
NObl10pt-Ob120pt		computation (of the bank reserve) (by its clerks)

We presented above the labels of subcategorization patterns as they appear at acquisition time. At processing time, there is no difference between Obl1 and Obl2, which are both of type Oblique. Our machine-guided approach helps the acquirer to select a rule as it only presents the relevant ones for a specific semantic type. For instance, in the case of a lexeme mapped to an OBJECT no rules having obliques will be presented to the acquirer as described below

example	semantics	subcat	lexical class
weapon	Obj	N	ObjN
father	Prop	NObl10pt	PropNOblEventObl10pt
bombing	Event	NObl10bl20pt	EventNOblObjObl10bj
computation		NObl10ptOb120pt	-EventObl120pt

The table above should be read as follows: the first column provides type examples for nouns, the second column (semantics) provides the list of semantic types that a noun can be, Obj (Object), Prop (Property) and Event, the third column (subcategorization) presents all subcategorizations a noun can subcategorize for, the fourth column (lexical class) concatenates the semantics and the subcategorization. For instance EventNOblObjObl1ObjEventObl120pt is the lexical class of nouns which are of type 'Event' and therefore subcategorize for two obliques (Obl) - the former must be Obj whereas the latter can be either Obj or Event. These Obl can be optional (Opt).

Acquirers may specify the preposition (head of the oblique or prepositional phrase). For instance, in the case of *father*, once an acquirer has mapped the word to the concept 'Father' which is a Prop (Property) the acquisition tool presents the subcategorization NObl10pt. This allows the acquirer to select which preposition(s) can go with the range of Father (in this case "of" will be selected). This information is important in generation. For generation, one must specify, at acquisition time, whether or not one can say *the bombing of Iraq*, *the bombing of Iraq by the US*, *the bombing by the US*. It also helps in word sense disambiguation.

In the case of verbs, one can also define lexico-syntactic classes for different semantic classes. For instance, in the case of ASSERTIVEACT the lexemes mapped to it will accept a comp clause (e.g. *he*

said (that) he would come) One class of aspectuals subcategorizes for nps (e.g. *I started a new book*), xcomps (e.g. *I started reading/writing a new book*), and accepts the intransitive alternation when the grammatical object is of type Event (e.g. *the surgery started very late*)<sup>3</sup>

## 4 Propagation of Lexicons

In this section, we briefly discuss how to extend a lexicon using derivational morphology, and off the shelf resources such as WordNet (Miller, 1990) to propagate the English lexicon with synonyms, and Levin's database of subcategorizations and alternations for English verbs (Levin, 1993) to encode syntactic information in the verb entries<sup>4</sup>

### 4.1 Morpho-semantics for Derivational Morphology

We refer the reader to Viegas *et al.* (1996) for the details on this type of acquisition and theoretical background of Lexical Rules (LRs). To sketch this operation briefly, applying morpho semantic LRs to the entry for the Spanish verb *comprar* (buy), our acquisition system produced automatically 26 new entries (*comprador-N1* (buyer), *comprable-Adj* (buyable), etc). This includes creating new syntax, semantics and syntax-semantic mappings with correct subcategorizations and also the right semantics. For instance, the lexical entry for *comprable* will have the subcategorization for predicative and attributive adjectives and the semantics adds the attribute FEASIBILITYATTRIBUTE to the basic meaning BUY of *comprar* (Viegas *et al.*, 1996) describes about 100 morpho-semantic LRs, which were applied to 1056 verb citation forms with 1,263 senses among them. The rules helped acquire an average of 26 candidate new entries per verb sense. This produced a total of 31,680 candidate entries, with an average of over 90% and 85% correctness in the assignment of syntax and semantics respectively. LRs constitute a powerful tool to extend a core lexicon from a monolingual viewpoint. We present other ways of extending lexicons, from monolingual (next paragraph) and multilingual (Section 5) perspectives.

### 4.2 Using WordNet

WordNet has been used as follows. We extracted the synsets associated with a lexeme using fuzzy string matches between, on the one hand, the value of the ontological concept (e.g., DESIRE), its definition (e.g., for DESIRE "to want something") and the concept and definition of its corresponding ISA concept (e.g., INTEND) and, on the other hand, the direct hypernyms and hyponyms for the lexeme in

<sup>3</sup>See Viegas *et al.* (1999) for the details on the web interfaces used for the acquisition.

<sup>4</sup>See Viegas *et al.* (1998) for more details.

WordNet synsets. For instance, for the English verb *expect*, mapped to the ontological concept DESIRE, our algorithm only kept one synset *hope, expect, trust, desire* for *expect* and the following synsets for its hypernyms *wish, desire, want*.

The output of our automatic procedure and manual filtering is illustrated below for the ontological concept DESIRE, along with the synonyms from WordNet belonging to the same ontological class.

DESIRE  
want  
expect  
trust  
wish

All these lexemes will be mapped to the concept DESIRE and minimally accept the same subcategorizations (e.g. np-v-np-xcomp as in *I want you to feel comfortable*).

We should mention that this step also involved some manual filtering by acquirers. We used a machine-guided mode to help the acquirer in this task. This type of filtering was done very quickly, mainly due to the fact that WordNet is organized on a semantic basis.

### 4.3 Using Levin's DB

One of the major problems in using Levin's database was filtering out homonyms, as classes in Levin's database are defined on the basis of the same subcategorization pattern (as seen in alternations) and not on a semantic basis, as shown by many researchers<sup>5</sup>.

The advantage of our approach is that it is semantic-based, this allows us to organize verbs into true (frame-based) semantic classes, with their associated sets of subcategorizations. Therefore, we can predict that all verbs belonging to a particular semantic class will have the same syntactic behavior. For instance, if one considers the semantic class of aspectual verbs which selects a theme of type Event, e.g. *begin, continue, finish*, then one can minimally associate to any verb belonging to this semantic class the following subcategorizations: (a) NP-V-NP in *John began his homework*, (b) NP-V-XCOMP *John began to work/working*. Note that the reverse is not necessarily true: verbs which accept (a) and (b) are not necessarily aspectuals, e.g. *forget* in *I forgot the key* or *I forgot to bring the key*.

## 5 Applicability to Other Languages

In this section, we briefly address what can be generalized to multiple languages. The methodologies described here are part of what is needed to build

<sup>5</sup>Many experiments have resulted in a similar finding, as described in (Dorr *et al.*, 1997), (Dang *et al.*, 1997). Saint-Dizier (1996) also showed that these classes do not apply easily to French.

a multi-purpose LKB while keeping the costs of acquisition as low as possible

### 5.1 Semantic Multilinguality

By mapping lexemes to concepts, it is possible to create lexicons for different languages, at a minimum cost, once a core lexicon has been acquired. This task can be further accelerated if one has access to bilingual dictionaries to semi-automate the translation task. Finally, if one has access to a rich structured ontology (as is the case in Mikrokosmos) then dynamic procedures (e.g., generalization, specialization) can help the acquirer in “filling” the gap in the case of lexico-semantic mismatches (e.g., *cook*, *bake* → *cuire*)

### 5.2 More Related Language

#### Multilinguality: Morpho-semantics

All the LRs (e.g., *LR2agent-of*) developed for Spanish can be used to extend other languages, even unrelated ones, in other words, these rules are language independent. The morpho-semantic aspect of the LRs is, however, specific to particular languages. But, in order to benefit from the work done on morpho-semantic LRs, we separated the assignment of affixes from the assignment of LRs. In other words, if in Spanish *LR2agent-of* is assigned to say the suffix *-dor*, by translating suffixes between languages, (*-dor* → *-eur* in French), the French lexicon can be extended in the same way (*comprador* → *acheteur*). Again, this work will necessitate some manual checking, because of some overgeneration, which cannot be accepted for generation. But overall, one can use the same methodology, the same LRs and engine to produce new entries.

### 5.3 Even More Related Language

#### Multilinguality

The subcategorizations attached to a lexeme have an even more idiosyncratic behavior than lexical LRs. But here again, the rules we developed can be applied at least to family-related languages, and then filtered out by a human. For instance, the Spanish word *comer* has the pattern *np-v-np* associated to it (e.g. *Juan come una pera*), so this same pattern will be attached to the translation of *comer* (eat) as in (*Juan eats a pear*). However, going from Spanish to English, one misses all the alternations (Levin, 1993) not common in Spanish such as *John gave Mary a book*

## 6 Summary

The Mikrokosmos lexicon acquisition group has acquired the following data - Spanish lexicon 7,000 word sense entries (35,000 word sense entries after applying the morpho-semantic lexical rules), Chinese lexicon about 3,000 word sense entries, and English, about 15,000 word sense entries so far. For

instance, the acquisition of 15,000 word sense entries took one year and involved 50% of the time of a computational linguist (to develop the methodology, train the acquirers and design the GUIs), 50 acquirer hours per week, 10 hours per week of a programmer to implement the GUIs, maintain the tools and test the entries.

Our approach to the development of lexicons differs from others in that our rules apply directly to semantic frames and not to the basic forms of verbs. Our methodology allows us to alleviate the burden of manual checking by applying linking rules directly on the semantics of the lexemes. Some rules add discourse related features, such as focus in some alternations, e.g., *they improved the situation* → *the situation improved*.

What is important to evaluate is how much do we gain by using rules and other resources. Today, it is still difficult to say exactly how much.

Adequately predicting the subcategorizations for a semantic class depends on its grain size: *the finer-grained, the better the prediction will be*. However, in NLP applications, where one is constrained by time, only the semantics necessary for a particular application is acquired, which means that in many cases the semantics is left at a coarser grain size than the one required to predict the subcategorizations. In practice, we overgenerate some subcategorizations and need therefore to have them checked by humans. This is why we have concentrated on a small set of rules. Results on that trade-off issue have been reported in Viegas *et al.* (1998).

Our experience in large-scale acquisition of lexicons shows that idiosyncrasies overrule many of our general rules. This is mainly due to the fact that we need a more fine-grained semantics than the one which is available now. This is not just a criticism of our framework, it is a general fact that we all encounter when investigating lexical semantics. This might be due to the fact that we work in a synchronic perspective (a highly recommended approach!), whereas language evolves constantly, thus creating “artificial” idiosyncrasies. In any case one cannot avoid them when building a computational semantic lexicon.

We have also learnt during the acquisition of the Mikrokosmos lexicons that different acquirers, who have been through the same intensive training, will arrive at the same number of meanings for a word, in more than 90% of the cases. The meaning of a word might differ, for different trained acquirers, along ISA links. Corpora also influence the decision of the acquirers, and here too we have seen some human “inconsistencies” which we think could be “corrected” automatically, as discussed in the following section.


## 7 Perspectives

We are investigating the issue of taking into account inconsistent lexical descriptions between the lexicon acquirers by taking advantage of the semantic information encoded in the ontology

If we look at the following data and their subcategorizations

- (1) I fixed the meal - [NP1, NP2]
- (2) I fixed a sandwich for you - [NP1, NP2, PP1]
- (3) I fixed you a sandwich - [NP1, NP3, NP2]

where *fix* means PREPAREFOOD, then one must allow the analysis and generation of any of sentences (1), (2) and (3), whether there is a mapping of *fix*<sup>B</sup> onto the concepts A or B or C<sup>6</sup>




Looking at example (1), and in absence of examples (2) and (3) in the corpus, *fix* could easily be mapped into A or C, whereas with examples (2) and (3), and in absence of example (1) in the corpus, it could be easily mapped into A or B by the acquirers. We claim that this is of no importance as far as there are mechanisms to go from one to the other. This requires to have access to semantic information. The diagram above is a computational linguist construct and has no "reality" per se. B and C are constructs which provide for every semantic class the different semantic patterns that a particular semantic class accepts, such as the pattern CREATEINGESTBENEF requires 3 semantic arguments (AGENT, BENEFICIARY and THEME), whereas CREATEINGESTTHEME only requires 2 semantic arguments. In this case, this means that the BENEFICIARY is optional, a fact the acquirer "failed" to recognize.

This diagram can be further specified for a particular natural language, where the required arguments are mapped to syntactic arguments and where lexical rules for a particular language provide the link between the different semantic patterns for a semantic class. The diagram below is for English where

<sup>6</sup>We give the general diagram for CREATEINGEST events, as PREPAREFOOD is a subtype and will inherit all the properties of the semantic class CREATEINGEST

subcategorizations in between square brackets are associated to the lexical items mapped to A, B and C



The corpus can indeed influence the way a lexicon acquirer will do the mapping. So if a lexicon acquirer creates an unspecified entry (mapping *fix* on (A), as opposed to (B) or (C)), dynamic mechanisms such as specialization or generalization would enable the system to get to (B) and (C) from (A) and vice versa (to (A) from (B) or (C)). Moreover, if a lexicon acquirer decides to map to (B) instead of (C) or vice versa, then a lexical rule (LR) between (B) and (C) will enable the system to go from (B) to (C) and vice versa.

In other words, although there are three potentially different ways of writing the lexicon entry for *fix* for example sentences (1), (2) and (3), these different ways of encoding *fix* should remain a virtual difference at processing time. The system must encode mechanisms and rules to "interpret" and reconcile the different points of view of different acquirers. This enables the system to process sentences (1), (2) and (3) from any of the three potential lexicon entries.

We believe that an important issue in computational semantics is to study how lexicon entries could be dynamically changed to fit different linguistic contexts and different acquirers' analysis of the data. This is what we plan to investigate in our future research.

### Acknowledgments.

This work has been supported in part by DoD under contract number MDA-904-92-C-5189. We are grateful to the Mikrokosmos team, and in particular Stephen Beale, Kavi Mahesh, Sergei Nirenburg, Boyan Onyshkevych and Victor Raskin. Without their input and the many discussions with them it would not have been possible for the author to develop the methodologies described here and supervise the acquisition of the Mikrokosmos lexicons. However, the opinions expressed in this ar-

article are those of the author and do not necessarily reflect their opinions. We would also like to thank the anonymous reviewers for their useful comments. Last but not least, we would like to thank all the acquirers who developed the lexicons over the years: Oscar Cossio, Ron Dolan, Margarita Gonzales, Wanying Jin, Julie Lonergan, Jeff Longwell, Maya, Javier Ochoa, Arnim Ruelas.

Large-scale Semantic LKB to Suit an Intelligent Planner In Proc of the 7th ENLG Toulouse

## References

- Dale, R. & C. Mellish (1998) Towards Evaluation in Natural Language Generation In Proc of LRE, Granada, Spain
- Dang, H., Rosenzweig, J. & M. Palmer (1997) Associating Semantic Components with Intersective Levin Classes In MCCS-97-314, CRL, NMSU
- Dorr, B., Olsen, M. & D. Clark (1997) Using WordNet to Posit Hierarchical Structure in Levin's Verb Classes In MCCS-97-314, CRL, NMSU
- Hovy, E. (1988) *Generating Natural Language Under Pragmatic Constraints* L. Erlbaum, Hillsdale, NJ
- Kilgarriff, A. (1997) *Sample the Lexicon* Technical Report, ITRI-97-01
- Levin, B. (1993) *English Verb Classes and Alternations: A Preliminary Investigation* Chicago University of Chicago Press
- Mahesh, K. (1996) *Ontology Development: Ideology and Methodology* MCCS-96-292, CRL, NMSU
- Miller, G. A. (1990) Nouns in WordNet: A Lexical Inheritance System In *International Journal of Lexicography* 3:4
- Nirenburg, S., J. Carbonell, M. Tomita & K. Goodman (1992) *Machine Translation: A Knowledge-Based Approach* Morgan Kaufmann
- Nirenburg, S. & C. DeFrise (1991) Practical Computational Linguistics In Johnson & Rosner (eds) *Computational Linguistics and Formal Semantics* CUP
- Saint-Dizier, P. (1996) Verb semantic classes based on 'alternations' and on WordNet-like semantic criteria: A Powerful Convergence In Proc of *Predicative Forms in Lexical Semantics and LKBs*, Toulouse
- Viegas, E., Onyshkevych, B., Raskin, V. & S. Nirenburg (1996) From *Submit* to *Submitted* via *Submission*: on Lexical Rules in Large-scale Lexicon Acquisition In Proc of the 34th ACL, UCSC, CA
- Viegas, E. (1999) The Manifesto of Large-scale Semantic Lexicon Acquisition In R. Zajac (ed) *TALN - Le Multilinguisme*, Vol. 40-1
- Viegas, E., Ruelas, A., Beale, S. & S. Nirenburg (1998) Extending a Core Lexicon Using On-line Language Resources with Savoir-Faire In Proc of the 1st LRE, Granada, Spain
- Viegas, E., Ruelas, A., Lonergan, J., Longwell, J., Beale, S. and S. Nirenburg (1999) Developing a

