Multilingual Computational Semantic Lexicons in Action: The WYSINNWYG Approach to NLP

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

Abstract

Much effort has been put into computational lexicons over the years, and most systems give much room to (lexical) semantic data. However, in these systems, the effort put on the study and representation of lexical items to express the underlying continuum existing in 1) language vagueness and polysemy, and 2) language gaps and mismatches, has remained embryonic. A sense enumeration approach fails from a theoretical point of view to capture the core meaning of words, let alone relate word meanings to one another, and complicates the task of NLP by multiplying ambiguities in analysis and choices in generation. In this paper, I study computational semantic lexicon representation from a multilingual point of view, reconciling different approaches to lexicon representation: i) vagueness for lexemes which have a more or less finer grained semantics with respect to other languages; ii) underspecification for lexemes which have multiple related facets; and, iii) lexical rules to relate systematic polysemy to systematic ambiguity. I build on a What You See Is Not Necessarily What You Get (WYSINNWYG) approach to provide the NLP system with the “right” lexical data already tuned towards a particular task. In order to do so, I argue for a lexical semantic approach to lexicon representation. I exemplify my study through a cross-linguistic investigation on spatially-based expressions.

1 A Cross-linguistic Investigation on Spatially-based Expressions

In this paper, I argue for computational semantic lexicons as **active knowledge sources** in order to provide Natural Language Processing (NLP) systems with the “right” lexical semantic representation to accomplish a particular task. In other words, lexicon entries are “pre-digested”, via a lexical processor, to best fit an NLP task. This What You See (in your lexicon) Is Not Necessarily What You Get (as input to your program) (WYSINNWYG) approach requires the adoption of a symbolic paradigm. Formally, I use a combination of three different approaches to lexicon represen-

tations: (1) lexico-semantic vagueness, for lexemes which have a more or less finer grained semantics with respect to other languages (for instance *en* in Spanish is vague between the Contact and Container senses of the Location, whereas in English it is finer grained, with *on* for the former and *in* for the latter); (2) lexico-semantic underspecification, for lexemes which have multiple related facets (such as for instance, *door* which is underspecified with respect to its Aperture or PhysicalObject meanings); and, (3) lexical rules, to relate systematic polysemy to systematic ambiguity (such as the Food Or Animal rule for *lamb*).

I illustrate the WYSINNWYG approach via a cross-linguistic investigation (English, French, Spanish) on spatially-based expressions, as lexicalised, for instance, in the prepositions *in*, *above*, *on*, ..., verbs *traverser*, (“go” across) in French, predicative nouns *montée*, (going up) in French, or in adjectives *upright*. Processing spatially-based expressions in a multilingual environment is a difficult problem as these lexemes exhibit a high degree of polysemy (in particular for prepositions) and of language gaps (i.e., when there is not a one-to-one mapping between languages, whatever the linguistic level; lexical, semantic, syntactic, etc). Therefore, processing these expressions or words in a multilingual environment minimally involves having a solution for treating: (a) syntactic divergences, *swim across* → *traverser ... à la nage* in French (cross ... swimming); (b) semantic mismatches, *river* translates into *fleuve*, *rivière* in French; and (c), cases which lie in between clear-cut cases of language gaps (*stand* → *se tenir debout/se tenir*, *lie* → *se tenir allongé/se tenir*). Researchers have dealt with a) and/or b), whereas WYSINNWYG presents a uniform treatment of a), b) and c), by allowing words to have their meanings vary in context.

In this paper, I restrict my cross-linguistic study to the (lexical) semantics of words with a focus on spatially-based expressions, and consider literal or non-figurative meanings only. In the next sections, I address representational problems which must be solved in order to best capture the phenom-

ena of ambiguity, polysemy and language gaps from a lexical semantic viewpoint. I then present three different ways of capturing the phenomena: lexico-semantic vagueness, lexico-semantic underspecification and lexical rules.

1.1 The Language Gap Problem

Upon a close examination of empirical data, it is often difficult to classify a translation pair as a syntactic divergence (e.g., Dorr, 1990; Levin and Nirenburg, 1993), as in *he limped up the stairs* → *il monta les marches en boitant* (French) (he went up the stairs limping) or a semantic mismatch (e.g., Palmer and Zhibiao, 1995; Kameyama et al., 1991), as in *lie, stand* → *se tenir* (French). Moreover, *lie* and *stand* could be translated as *se tenir couché/allongé* (be lying) and *se tenir debout* (be up) respectively, thus presenting a case of divergence, or they could both be translated into French as *se tenir*, thus presenting a case of conflation, (Talmy, 1985). Depending on the semantics of the first argument, one might want to generate the divergence, (e.g., *se tenir debout/couché*), or not (e.g., *se tenir*), thus considering *se tenir* as a mismatch as in (1):

- (1) *Pablo se tenait au milieu de la chambre.*
(Sartre)
(Pablo stood in the middle of the bedroom.)

In order to account for all these language variations, one cannot “freeze” the meanings of language pairs. In section 2.1, I show that by adopting a continuum perspective, that is using a knowledge-based approach where I make the distinction between lexical and semantic knowledge, cases in between syntactic divergences and semantic mismatches (*se tenir*) can be accounted for in a uniform way. Practically, the proposed method can be applied to interlingua approaches and transfer approaches, when these latter encode a layer of semantic information.

1.2 The Lexicon Representation Problem

Within the paradigm of knowledge-based approaches, there are still lexicon representation issues to be addressed in order to treat these language gaps. It has been well documented in the literature of this past decade that a sense enumeration approach fails from a theoretical point of view to capture the core meaning of words (e.g., (Ostler and Atkins, 1992), (Boguraev and Pustejovsky, 1990),...) and complicates from a practical viewpoint the task of NLP by multiplying ambiguities in analysis and choices in generation.

Within Machine Translation (MT), this approach has led researchers to “add” ambiguity in a language which did not have it from a monolingual perspective. Ambiguity is added at the lexical level within transfer based approaches (“river1” → “rivière”; “river2” → “fleuve”); and at the semantic

level within interlingua based approaches (“rivière” → RIVER - DESTINATION: RIVER; “fleuve” → RIVER - DESTINATION: SEA; “river” → RIVER - DESTINATION: SEA, RIVER), whereas again “river” in English is not ambiguous with respect to its destination.

In this paper, I show that ambiguity can be minimised if one stops considering knowledge sources as “static” ones in order to consider them as **active** ones instead. More specifically, I show that building on a computational theory of lexico-semantic vagueness and underspecification which merges computational concerns with theoretical concerns enables an NLP system to cope with polysemy and language gaps in a more effective way.

Let us consider the following simplified input semantics (IS):

- (2) PositionState(Theme:Plate,Location:Table),

This can be generated in Spanish as *El plato esta en la mesa*; where Location is lexicalised as *en* in Figure 1.

To generate (2) into English, requires the system to further specify Location for English as LocContact, in order to generate *The plate is on the table*, where *on* corresponds to the Spanish *en1*, sub-sense of *en*, as shown in Figure 1.

Figure 1: Subset of the Semantic Types for Prepositions

From a monolingual perspective, there is no need to differentiate in Spanish between the 3 types of Location as LocContact, LocContainer and LocBuilding, as these distinctions are irrelevant for Span-

ish analysis or generation, with respect to Figure 1. However, within a multilingual framework, it becomes necessary to further distinguish Location, in order to generate English from (2). In the next sections, I will show that lexical semantic hierarchies are better suited to account for polysemous lexemes than lexical or semantic hierarchies alone, for multilingual (and monolingual) processing.

2 The WYSINNWYG Approach

I argue that treating lexical ambiguity or polysemy and language gaps computationally requires 1) fine-grained lexical semantic type hierarchies, and 2) to allow words to have their meanings vary in context.

Much effort has been put into lexicons over the years, and most systems give more room to lexical data. However, most approaches to lexicon representation in NLP systems have been motivated more by computational concerns (economy, efficiency) than by the desire for a computational linguistic account, where the concern of explaining a phenomenon is as important as pure computational concerns. In this paper, I adopt a computational linguistic perspective, showing however, how these representations are best fitted to serve knowledge-driven NLP systems.

2.1 A Continuum Perspective on Language Gaps

I argue that resolving language gaps (divergences, mismatches, and cases in between) is a generation issue and minimally involves:

- 1) using a knowledge-based approach to represent the lexical semantics of lexemes;
- 2) developing a computational theory of lexico-semantic vagueness, underspecification, and lexical rules;

In this paper, I only address lexical representational issues, leaving the generation issues (such as the use of planning techniques, the integration of the process in lexical choice) aside.¹

I illustrate through some examples below, how a compositional semantics approach, e.g. knowledge-based, can help in dealing with language gaps.² I will use the French (*se tenir*) and English (*stand*, *lie*) simplified entries below, in my illustration of mismatches between the generator and the lexicons. Semantic types are coded in the sense feature:

¹Generation issues are fully discussed in (Beale and Viegas, 1996). This first implementation of some language gaps has a very limited capability for the treatment of vagueness and underspecification; although it takes advantage of the semantic type hierarchy, it still lacks the benefit of having the lexical type hierarchy presented here.

²Note that absence of compositionality, such as in idioms *kick the (proverbial) bucket* or syntagmatic expressions *heavy smoker*, is coded in the lexicon.

```
[key: "se-tenir3",
 form: [orth: [ exp: "se tenir"]],
 sense: [sem: [name: Position-state], ...]
```

```
[key: "stand2",
 form: [orth: [ exp: "stand"]],
 sense: [sem: [name: PsVertical], ...]
```

```
[key: "lie1",
 form: [orth: [ exp: "lie"]],
 sense: [sem: [name: PsHorizontal], ...]
```

Figure 2 illustrates a subset of the Semantic Type Hierarchy (STH) common to all dictionaries and of two subsets of the Lexical Type Hierarchy (LTH) for French and English.

Figure 2: Example of an STH linked to a Fragment of the French and English LTHs.

I illustrate below three main types of gaps between the input semantics (IS) to the generator and the lexicon entries (LEX) of the language in which to generate. I focus on the generation of the predicate:

(i) **IS - LEX exact match** Generating, in French, from the simplified IS below (3),

(3) **PositionState(agent: john, against: wall)**

is easy as there is a single French word in (3) that lexicalises the concept PositionState, which is *se tenir*. Therefore *se tenir* is generated in *John se tenait contre le mur* (John was/(stood) against the wall).

(ii) **IS - LEX vagueness** Generating, in French, from the partial IS below (4),

(4) `PsVertical(agent:john,against:wall)`

needs extra work from the generator, with respect to the lexicon entry for French. In Figure 2, one can see in STH that `PsVertical` is a sub-type of `PositionState`, which has a mapping in LTH for French to *se-tenir*³. This illustrates a case of vagueness between English and French. In this case, the generator will generate the same sentence *John se tenait contre le mur*, as is the case for the exact match in (i). Note that generating the divergence *se tenait debout* (stand upright) although correct and grammatical, would emphasise the position of *John* which was not necessarily focused in (4). The divergence can be generated by “composing” `PsVertical` as `PositionState` (lexicalised as *se tenir*) and `Vertical` (lexicalised as *debout*).

(iii) **IS - LEX Underspecification** Generating, in French, from the partial IS below (5),

(5) `PsVertical(agent:john,against:wall, time:t1) & PsHorizontal(agent:john, against:wall,time:t2) & t1<t2`

needs extra work from the lexicon processor, with respect to the entries presented here, as one does not want to end up generating *John se tint contre le mur puis il se tint contre le mur* (John was against the wall then he was against the wall). Because of the conjunctions here, one cannot just consider *se tenir* as vague with respect to *lie* and *stand*. This illustrates a lexicon in action, where the lexical processor must process *se tenir* as underspecified:

`PositionState` $\rightarrow$ `PsVertical` $\vee$ `PsHorizontal`

The lexical processor will thus produce the divergences *se tenir debout* (stand) and *se tenir allongé* (lying) to generate (with some generation processing such as lexical choice, ellipsis, pronominalisation, etc) *John se tenait (debout) contre le mur puis s'allongea contre lui* (John was standing against the wall then he lied against it).

Where the continuum perspective comes in, is that we do not want to “freeze” the meanings of words once and for all. As we just saw, in French one might want to generate *se tenir debout* or just *se tenir* depending on the semantics of its arguments and also depending on the context as in (5).

In the WYSINWYG approach, words are allowed to have their “meanings” vary in context. In other words, the literal meaning(s) coded in the lexicon is/are the “closest” possible meaning(s) of a word within the STH context, and by enriching the discourse context (dc), one ends up “specialising”

or “generalising” the meaning(s) of the word, using formally two hierarchies: semantic (STH) and lexical (LTH), enabling different types of lexicon representations: vagueness, underspecification and lexical rules.

2.2 A Truly Multilingual Hierarchy

Multilingual lexicons are usually monolingual lexicons connected via translation links (Tlinks), whereas truly multilingual lexicons, as defined by (Cahill and Gazdar, 1995), involve $n + 1$ hierarchies, thus involving an additional abstract hierarchy containing information shared by two or more languages. Figure 3 illustrates the STH which is shared by all lexicons (French, English, Spanish, etc), and the lexical MLTH which involves the abstract hierarchy shared by all LTHs.

Figure 3: Subset of the Multilingual Hierarchy for Prepositions

The lexicons themselves are also organised as language lexical type hierarchies (Spanish LTH, English LTH in Figure 3). For instance, the English dictionary (*eng-lexeme*) has the English prepositions (*eng-prep*) as one of its sub-types, which itself has as sub-types all the English prepositions (*along*, *through*, *on*, *in*, ...). These prepositions have in turn subtypes (for instance, *on* has *on1*, *on2*, ...), which can themselves have subtypes (*on11*, *on12*, ...). All these language dependent LTHs inherit part of their information from a truly Multilingual Lexical Type Hi-

erarchy (MLTH), which contains information shared by all lexicons. There might be several levels of sharing, for instance, family-related languages sharing.

Lexical types are linked to the STH via their language LTH and the MLTH, so that these lexicons can be used by either monolingual or multilingual processing. The advantages of a MTLH extend to 1) lexicon acquisition, by allowing lexicons to inherit information from the abstract level hierarchy. This is even more useful when acquiring family-related languages; and 2) robustness, as the lexical processors can try to “make guesses” on the assignment of a sense to a lexeme absent from a dictionary, based on similarities in morphology or orthography, with other family-related language lexemes.³

2.3 Vagueness, Underspecification and Lexical Rules

The STH along with the LTH allow the lexicographers to leave the meaning of some lexemes as vague or underspecified. The vagueness or underspecification typing allows the lexical processor to specialise or generalise the meaning of a lexeme, for a particular task and on a needed basis. Formally, generalisation and specialisation can be done in various ways, as specified for instance in (Kameyama et al., 1991), (Poesio, 1996), (Mahesh et al., 1997).

2.3.1 Lexicon Vagueness

A lexicon entry is considered as vague when its semantics is typed using a general monomorphic type covering multiple senses, as is the case of the French entry “se-ténir³”, or the Spanish preposition *en*, as represented in (6).

(6) [key: “en”,
form: [orth: [exp: “en”], ...
sense: [sem: [name: Location], ...]

It is at processing time, and only if needed, that the semantic type Location for *en* can be further processed as LocContact, LocContainer, ... to generate the English prepositions (on, at, ...).

Lexicon vagueness is represented by mapping the citation form *lex* of any word *x* appearing in a corpus to a semantic monomorphic type *m*, which belongs to STH. Let us consider MAPS, the function which links *lex* to STH, *dc* a discourse context where *lex* can appear, and $\sqsubseteq$ the immediate type/sub-type relation between types of STH, then:

(7) *x* is vague iff
 $\exists m \in STH : m = MAPS(dc, lex(x)) \wedge$
 $\exists n, o \in STH : n \sqsubseteq m \wedge o \sqsubseteq m \wedge n \neq o \wedge$
 $\forall r \in STH : r \sqsubseteq m : \exists q \in STH : q \sqsubseteq r$

³I have not investigated this issue yet, but see (Cahill, 1998) for promising results with respect to making guesses on phonology.

In other words, *lex* is vague, if *m* is in a type/sub-type relation with all its immediate sub-types.

2.3.2 Lexicon Underspecification

The meaning of a lexeme is considered as underspecified when its semantics is represented via a polymorphic type, which presents a disjunction of semantic types,⁴ thus covering different **polysemous** senses, as is the case of the Spanish preposition “por” in (8), and typical examples in lexical semantics, such as *door* which is typed as PHYSICAL_OBJECT-OR-APERTURE.⁵

(8) [key: “por”,
form: [orth: [exp: “por”], ...
sense: [sem: [name: Through; Along], ...]

It is at processing time only, and on a needed basis only, that the semantic type Through-OR-Along for *por* can be further processed as either Through, or Along, ..., thus allowing the generator or analyser to find the appropriate representation depending on the task. Disambiguating “por” to generate English, requires that the lexeme be embedded within the discourse context, where the filled arguments of the prepositions will provide semantic information under constraints. For instance, *walk* and *river* could contribute to the disambiguation of *por* as Along.

Lexicon underspecification is represented by mapping *lex* (the citation form of a word *x*) to a semantic polymorphic type *p*, which belongs to STH, then:

(9) *x* is underspecified iff
 $\exists p \in STH : m = MAPS(dc, lex(x)) \wedge$
 $\exists s \subseteq STH : p = \vee s \wedge Card(s) \geq 2$

In other words, *lex* is underspecified, if *p* is a disjunction of types, and no type/sub-type relation is required.

⁴See (Sanfillippo, 1998) and (Buitelaar, 1997) for different computational treatments of underspecified representations. The former deals with multiple subcategorisations (whereas I am also interested in polysemous senses), the latter includes homonyms, which I agree with Pinkal (1995) should be left apart.

⁵I believe that lexico-semantic underspecification is concerned with polysemous lexemes only (such as *door*, *book*, etc) and not homonyms (such as *bank* as financial-bank or river-bank) called H-Type ambiguous in (Pinkal, 1995). I believe the H-Type ambiguous lexemes should be related via their lexical form only, while their semantic types should remain unrelated, i.e., there is no need to introduce a “disjunction fallacy” as in (Poesio, 1996). It might be the case that homonyms require pragmatic underspecification as suggested, for instance, in (Nunberg, 1979), but in any case are beyond the scope of this paper.

2.4 Lexical Rules

Lexical rules (LRs) are used in WYSINNWYG to relate systematic ambiguity to systematic polysemy. They seem more appropriate than underspecification for relating the meanings of lexemes such as “lamb” or “haddock” which can be either of type Animal or Food (Pustejovsky, 1995, pp. 224). LRs and their application time in NLP have received a lot of attention (e.g., Copestake and Briscoe, 1996; Viegas et al., 1996), therefore, I will not develop them further in this paper, as the rules themselves activated by the lexical processor produce different entries, with neither type/sub-type relations nor disjunction between the semantic types of the old and new entries. In WYSINNWYG, lexicon entries related via LRs are neither vague nor underspecified. For instance, the “grinding rule” of Copestake and Briscoe for linking the systematic Animal - Food polysemy as in *mutton* / *sheep* or in French where we have a conflation in *mouton*, allows us to link the entries in English and sub-senses in French, without having to cope with the semantic “disjunction fallacy problem” of (Poesio, 1996).

3 Conclusions - Perspectives

I have argued for **active knowledge sources** within a knowledge-based approach, so that lexicon entries can be processed to best fit a particular NLP task. I adopted a computational linguistic perspective in order to explain language phenomena such as language gaps and polysemy. I argued for semantic and lexical type hierarchies. The former is shared by all dictionaries, whereas the latter can be organised as a truly multilingual hierarchy. In that respect, this work differs from (Han et al., 1996) in that I do not suggest an ontology per language, but argue on the contrary for one semantic hierarchy shared by all dictionaries.⁶ Other works which have dealt with mismatches, e.g., (Dorr and Voss, 1998) with their interlingua and knowledge representations, (Sérasset, 1994) with his “interlingua acceptations”, or (Kameyama, et al, 1991) with their infons, cannot account for cases which lie in between clear-cut cases of divergences and mismatches such as the example “se tenir” discussed in this paper.

I have shown that enabling lexicon entries to be typed as either lexically vague or underspecified, or linked via LRs, allows us to account for the variations of word meanings in different discourse contexts. Most of the works in computational lexical semantics have dealt with either underspecification or LRs, trying to favour one representation over the other. There was previously no computational treat-

ment of lexical semantic vagueness. In discourse approaches and formal semantics, the use of underspecification in terms of truth values led researchers, when applying their research to individual words, to the “disjunction fallacy problem”, where a person who went to the bank, ended up going to the (financial-institution OR river-shore), whatever this object might be!, instead of a) going to the financial-institution OR b) going to the river-shore.

In this paper, I have presented the usefulness of each representation, depending on the phenomenon covered. I showed the need to consider underspecification for polysemous items only, leaving homonyms to be related via their lexical forms only (and not their semantics). I believe that LRs have room for polysemous lexemes such as the *lamb* example, as here again one could not possibly imagine an animal being (food-OR-animal) in the same discourse context.⁷

Finally, lexical vagueness enables a system to process lexical items from a multilingual viewpoint, when a lexeme becomes ambiguous with respect to another language. From a multilingual perspective, there is no need to address the “sororites paradox” (Williamson, 1994), which tries to put a clear-cut between values of the same word (e.g., *not tall ... tall*). It is important to note that WYSINNWYG accepts redundancy in the lexicon representations: lexemes can be both vague and underspecified or either one.

One could object that the WYSINNWYG approach is knowledge intensive and puts the burden on the lexicon, as it requires one to build several type hierarchies: a STH shared by all languages and a LTH per language which inherits from the MLTH. However, the advantages of the WYSINNWYG approach are many. First, by using the MLTH, acquisition costs can be minimised, as a lot of information can be inherited by lexicons of family-related languages. This multilingual approach has been successfully applied to phonology by (Cahill and Gazdar, 1995). Second, the task of determining the meaning of words requires human intervention, and thus involves some subjectivity. WYSINNWYG presents a good way of “reconciling” different lexicographers’ viewpoints by allowing a lexical processor to specialise or generalise meanings on needed basis. As such, whether a lexicographer decides to sense-tag “en” as Location or creates the sub-senses “en1” and “en2” remains a virtual difference for the NLP system. Finally, and most important, WYSINNWYG presents a typing environment which accounts for the flexibility of word meanings in context, thus allowing lexicon acquirers to map words to their “closest” core meaning within STH (e.g., “se

⁶However, I do not preclude that there might be different views on the semantic hierarchy depending on the languages considered: “filters” could be applied to the STH to only show the relevant parts of it for some family-related languages.

⁷The fact that some cultures eat “living” creatures would require to type these lexemes using underspecification (food-OR-animal) instead of a lexical rule in their cultures.

tenir" → PositionState) and use mechanisms (such as generalisation, specialisation) to modulate their meanings in context (e.g., "se tenir" → PsVertical). In other words, WYSINWYG helps not only in sense selection but also in sense modulation.

Further research involves investigating representation formalisms, as discussed in (Briscoe et al., 1993) to best implement these type inheritance hierarchies.

4 Acknowledgements

This work has been supported in part by DoD under contract number MDA-904-92-C-5189. I would like to thank my colleagues at CRL for commenting on a former version of this paper. I am also grateful to John Barnden, Pierrette Bouillon, Boyan Onyshkevych, Martha Palmer, and the anonymous reviewers for their useful comments.

References

- S. Beale and E. Viegas. 1996. Intelligent Planning meets Intelligent Planners. In *Proceedings of the Workshop on Gaps and Bridges: New Directions in Planning and Natural Language Generation, at ECAI'96*, Budapest, 59-64.
- B. Boguraev and J. Pustejovsky. 1990. *Knowledge Representation and Acquisition from Dictionary*. Coling Tutorial, Helsinki, Finland.
- T. Briscoe, V. de Paiva and A. Copestake (eds). 1993. *Inheritance, Defaults, and the Lexicon*. Cambridge University Press.
- P. Buitelaar. 1997. A Lexicon for Underspecified Semantic Tagging. In *Proceedings of the Siglex Workshop on Tagging Text with Lexical Semantics: Why, What, and How?*, Washington DC.
- L. Cahill and G. Gazdar. 1995. Multilingual Lexicons for Related Lexicons. In *Proceedings of the 2nd DTI Language Engineering Conference*.
- L. Cahill. 1998. Automatic extension of a hierarchical multilingual lexicon. In *Proceedings of the Second Multilinguality in the Lexicon Workshop*, sponsored by the 13th biennial European Conference on Artificial Intelligence (ECAI-98).
- A. Copestake and T. Briscoe. 1996. Semi-Productive Polysemy and Sense Extension. In *Journal of Semantics*, vol.12.
- B. Dorr. 1990. Solving Thematic Divergences in Machine Translation. In *Proceedings of the 28th Annual Meeting of the Association for Computational Linguistics*.
- C. Han, F. Xia, M. Palmer, J. Rosenzweig. 1996. Capturing Language Specific Constraints on Lexical Selection with Feature-Based Lexicalized Tree-Adjoining Grammars. in *Proceedings of the International Conference on Chinese Computing* Singapore.
- M. Kameyama, R. Ochitani and S. Peters. 1991. Resolving Translation Mismatches With Information Flow. In *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*.
- R. Keefe and P. Smith. (eds) 1996. *Vagueness: a Reader*. A Bradford Book. The MIT Press.
- L. Levin and S. Nirenburg. 1993. Principles and Idiosyncrasies in MT Lexicons, In *Proceedings of the 1993 Spring Symposium on Building Lexicons for Machine Translation*, Stanford, CA.
- K. Mahesh, S. Nirenburg and S. Beale. 1997. If You Have It, Flaunt It: Using Full Ontological Knowledge for Word Sense Disambiguation. In *Proceedings of the 7th International Conference on Theoretical and Methodological Issues in Machine Translation*.
- G. Nunberg. 1979. The Non-uniqueness of Semantic Solutions: Polysemy. *Linguistics and Philosophy* 3.
- N. Ostler and S. Atkins. 1992. Predictable meaning shift: Some linguistic properties of lexical implication rules. In Pustejovsky and Bergler (eds.) *Lexical Semantics and Knowledge Representation*. Springer Verlag.
- M. Palmer and W. Zhibiao. 1995. Verb Semantics for English-Chinese Translation. *Machine Translation*, Volume 10, Nos 1-2.
- M. Pinkal. 1995. *Logic and Lexicon*. Oxford.
- M. Poesio. 1996. Semantic Ambiguity and Perceived Ambiguity. In K. van Deemter and S. Peters (eds.) *Semantic Ambiguity and Underspecification*.
- J. Pustejovsky. 1995. *The Generative Lexicon*. MIT Press.
- A. Sanfilippo. 1998. Lexical Underspecification and Word Disambiguation. In E. Viegas (ed.) *Breadth and Depth of Semantic Lexicons*. Kluwer Academic Press.
- G. Sérasset. 1994. *SUBLIM: un système universel de bases lexicales multilingues et NADIA: sa spécialisation aux bases lexicales interlingues par acceptions*. PhD. Thesis, GETA, Université de Grenoble.
- L. Talmy. 1985. Lexicalization Patterns: semantic structure in lexical forms. In Shopen (ed), *Language Typology and Syntactic Description III*. CUP.
- E. Viegas, B. Onyshkevych, V. Raskin and S. Nirenburg. 1996. From *Submit* to *Submitted* via *Submission*: on Lexical Rules in Large-scale Lexicon Acquisition. In *Proceedings of the 34th Annual meeting of the Association for Computational Linguistics*, CA.
- C. Voss and B. Dorr. 1998. Lexical Allocation in IL-Based MT of Spatial Expressions. In P. Olivier and K.-P. Gapp (eds.) *Representation and Processing of Spatial Expressions*. Lawrence Erlbaum Associates.
- T. Williamson. 1994. *Vagueness*. Routledge.