

Semi-automatic Induction of Systematic Polysemy from WordNet

Noriko Tomuro
DePaul University

School of Computer Science, Telecommunications and Information Systems
243 S. Wabash Ave.
Chicago IL 60604
cphdnt@ted.cs.depaul.edu

Abstract

This paper describes a semi-automatic method of inducing *underspecified semantic classes* from WordNet verbs and nouns. An underspecified semantic class is an abstract semantic class which encodes *systematic polysemy*: a set of word senses that are related in systematic and predictable ways. We show the usefulness of the induced classes in the semantic interpretations and contextual inferences of real-word texts by applying them to the predicate-argument structures in Brown corpus.

1 Introduction

WordNet (Miller, 1990) has been used as a general resource of broad-coverage lexical information in many Natural Language Processing (NLP) tasks, including sense tagging, text summarization and machine translation. However, like other large-scale knowledge-base systems or machine readable dictionaries (MRDs), WordNet contains massive ambiguity and redundancy. In particular, since WordNet senses are more fine-grained than most other MRDs such as LDOCE (Procter, 1978), each word entry is more ambiguous. For example, WordNet 1.6 (released December 1997) lists the following 9 senses for the verb *write*:

1. **write**, compose, pen, indite – produce a literary work
2. **write** – communicate or express by writing
3. publish, **write** – have (one's written work) issued for publication
4. **write**, drop a line – communicate (with) in writing
5. **write** – communicate by letter
6. compose, **write** – write music
7. **write** – mark or trace on a surface
8. **write** – record data on a computer
9. spell, **write** – write or name the letters

These fine sense distinctions may not be desired in some applications. Consequently any system which incorporates WordNet without customization must presume this redundancy, and may need to control the ambiguities in order to make the computation tractable.

Although the redundancy in WordNet could be a drawback, it can be an ideal resource for a broad-coverage domain-independent semantic lexicon based on *underspecified semantic classes* (Buitelaar, 1997, 1998). An underspecified semantic class is an abstract semantic type which encodes *systematic polysemy* (or *regular polysemy* (Apresjan, 1973)):¹ a set of word senses that are related in systematic and predictable ways (eg. INSTITUTION and BUILDING meanings of the word *school*). These related word senses are grouped together, and assigned an abstract semantic class that generalizes the relation. This way, we do not need to distinguish or disambiguate word senses that encompass several semantic "axes", and we can regard an underspecified class as a multi-dimensional semantic entity. This abstract class is underspecified because it does not specify either one of the member senses. Here, in building a lexicon based on such underspecified semantic classes, redundancy in WordNet is a desirable property since the amount of information lost by abstraction is minimized. Also, since WordNet sense entries are taken from general but wide range of domains, systematic polysemy can be extracted from the dictionary rather than from a sense-tagged corpus. Therefore, data sparseness problems become less significant. Then, the resulting lexicon can effectively compact the redundancy and ambiguity in WordNet by two dimensions: abstraction and systematic polysemy.

The use of underspecified semantic classes is one of the *underspecification* techniques being investigated in recent years (van Deemter and Peters,

¹Note that systematic polysemy should be contrasted with *homonymy* which refers to words which have more than one unrelated sense (eg. FINANCIAL_INSTITUTION and SLOPING_LAND meanings of the word *bank*).

1996). This underspecified class has several advantages. First, it can compactly represent the ambiguity which arises from multiple related senses. Thus it is more expressive and computationally efficient than single sense representations. Second, it can facilitate abductive inference through the systematicity between senses: given a word with n related senses, the identification of one sense in a context can imply maximally all n senses, some of which may only be implicit in the context. In addition, when two systematically polysemous words are used together, the combination enables even more powerful inferences through a complex matching between the two sets of systematic relations. Then, a domain-independent broad-coverage lexicon defined by such abstract underspecified classes can be used as a *background lexicon* in domain-specific reasoning tasks such as Information Extraction (Kilgarriff, 1997), or as a general semantic lexicon for parsing, as well as for many other NLP tasks that require contextual inferences.

However, automatic acquisition of systematic polysemy has been a difficult task. In fact, in most previous work in lexical semantics it is done manually (Buitelaar, 1997, 1998). In this paper, we present a semi-automatic method of inducing underspecified semantic classes from WordNet verbs and nouns. The method first applies a statistical analysis to obtain a rough approximation of the sense *dependencies* found in WordNet. Incorrect dependencies are then manually filtered out. Although the approach is not fully automated, it provides a principled way of acquiring systematic polysemy from a large-scale lexical resource, and greatly reduces the amount of manual effort that was previously required. Furthermore, by having a manual intervention, the results will be able to reflect our prior knowledge about WordNet that was not assumed in the statistical analysis. To see the usefulness of the induced semantic classes in the contextual inferences of real-world texts, predicate-argument structures are extracted from Brown corpus, and the occurrences of such classes are observed.

2 Systematic Polysemy

Before presenting the induction method, we first clarify what we consider a systematic polysemy in the work described in this paper, and explain the assumptions we made for such polysemy.

Our systematic polysemy is analogous to *logical polysemy* in (Pustejovsky, 1995): word senses in which there is no change in lexical category, and the multiple senses of the word have overlapping, dependent, or shared meanings. This definition excludes meanings obtained by cross-categorical alternations (eg. denominals) or morphological alternations (eg. suffixing with *-ify*), or homonyms or metaphors, and includes only the senses of the word of the same cat-

egory and form that have some systematic relations. For example, INSTITUTION and BUILDING meanings of the word *school* are systematically polysemous because BUILDING relates to INSTITUTION by the location of the institution.

For nouns, each polysemous sense often refers to a different object. In the above example, *school* as INSTITUTION refers to an organization, whereas *school* as BUILDING refers to a physical object. On the other hand, for verbs, polysemous senses refer to different aspects of the same action. For example, a word *write* in the sentence “John wrote the book” is ambiguous between CREATION (of the book) and COMMUNICATION (through the content of the book) meanings. But they both describe the same action of John writing the particular book. Here, these two meanings are systematically related by referring to the *causation* aspect (CREATION) or the *purpose* aspect (COMMUNICATION) of the *write* action. This view is largely consistent with the *entailment* relations (temporal inclusion and causation) used to organize WordNet verb taxonomies (Fellbaum, 1990).

Another assumption we made is the dependency between related senses. In the work in this paper, sense dependency is viewed as *sense extension*, similar to (Copestake and Briscoe, 1995), in which a primary sense causes the existence of secondary senses. This assumption is in accord with *lexical rules* (Copestake and Briscoe, 1995; Ostler and Atkins, 1992), where meaning extension is expressed by if-then implication rules. In the above example of the noun *school*, INSTITUTION meaning is considered as the primary and BUILDING as the secondary, since institutions are likely to have office space but building may be occupied by other entities besides institutions. Similarly for the verb *write*, CREATION is considered as the primary and COMMUNICATION as the secondary, since communication takes place through the object that is just produced but communication can take place without producing an object.

3 Induction Method

Our induction method is semi-automatic, requiring a manual filtering step between the phased automatic processing. The basic scheme of our method is to first identify the prominent pair-wise cooccurrence between any two basic types (abstract senses), and then build more complex types (underspecified classes) by the composition of those cooccurrences. But instead of generating/composing all possible types statically, we only maintain the pair-wise relations in a graph representation called *type dependency graph*, and dynamically form/induce the underspecified classes during the phase when each WordNet entry is assigned the class label(s).

Based on the definitions and assumptions described in the previous section 2, underspecified semantic classes are induced from WordNet 1.6 (released December 1997) by the following steps:

1. Select a set of abstract (coarse-grained) senses from WordNet taxonomies as basic semantic types. This step is done manually, to determine the right level of abstraction to capture systematic polysemy.
2. Create a *type dependency graph* from ambiguous words in WordNet. This step is done by two phased analyses: an automatic analysis followed by a manual filtering.
3. Generate a set of underspecified semantic classes by partitioning the senses of each word into a set of basic types. Each set becomes an underspecified semantic class. This step is fully automatic.

Each step is described in detail below.

3.1 Coarse-grained Basic Types

As has been pointed out previously, there are many regularities between polysemous senses, and these regularities seem to hold across words. For example, words such as *chicken* and *duck* which have ANIMAL sense often have MEAT meaning also (i.e., *animal-grinding* lexical rule (Copestake and Briscoe, 1992)). This generalization holds at an abstract level rather than the word sense level. Therefore, the first step in the induction is to select a set of abstract senses that are useful in capturing the systematicity. To this end, WordNet is a good resource because word senses (or *synsets*) are organized in taxonomies.

Ideally, basic types should be semantically orthogonal, to function essentially as the “axes” in a high-dimensional semantic space. Good candidates would be the top abstract nodes in the WordNet taxonomies or lexicographers’ file names listed in the sense entries. However, both of them fall short of forming a set of orthogonal axes because of several reasons. First, domain categories are mixed in with ontological categories (eg. *competition* and *body verb* categories). Second, some categories are ontologically more general than others (eg. *change* category in verbs). Third, particularly for the verbs, senses that seem to take different argument noun types are found under the same category (eg. “*ingest*” and “*use*” in *consumption* category). Therefore, some WordNet categories are broken into more specific types.

For the verbs, the following 18 abstract basic types are selected.

```
change(CH A) communication(COMM)
cognition(COG) competition(COMP)
contact(CONT) motion(MOT)
emotion(EMO) perception(PER)
possession(POSS) stative(STA)
weather(WEA) ingestion(ING)
use(USE) social(SOC) body(BOD)
phy_creation(PC R) mental_creation(MCR)
verbal_creation(VCR)
```

These are mostly taken from the classifications made by lexicographers. Two classes (“consumption” and “creation” are subdivided into finer categories (*ingestion*, *use* and *physical/mental/verbal_creation*, respectively) according to the different predicate-argument structures they take.

For the nouns, 31 basic types are selected from WordNet top categories (unique beginners):²

```
entity(ENT) life_form(LIF)
causal_agent(AGT) human(HUM)
animal(ANI) plant(PLA) object(OBJ)
natural_object(NOBJ) substance(SUB)
food(FOOD) artifact(AFT) article(ART)
location(LOC) psych_feature(PHY)
cognition(COG) feeling(FEEL)
motivation(MOT) abstraction(ABS)
time(TIME) space(SPA) attribute(ATT)
relation(REL) social_relation(SREL)
communication(COMM) shape(SHA)
measure(MEA) event(EVE) action(ACT)
possession(POSS) state(STA)
phenomena(PHE)
```

Senses under the lexicographers’ class “group” are redirected to other classes, assuming a collection of a type has the same basic semantic properties as the individual type.

3.2 Type Dependency Graph

After the basic types are selected, the next step is to create a *type dependency graph*: a directed graph in which nodes represent the basic types, and directed edges correspond to the systematic relations between two basic types.

The type dependency graph is constructed by an automatic statistical analysis followed by a manual filtering process, as described below. The premise here is that, if there is a systematic relation between two types, and if the regularity is prominent, it can be captured by the type cooccurrence statistics. In machine learning, several statistical techniques have been developed which discover dependencies among features (or causal structures), such

²Noun top categories in WordNet do not match exactly with lexicographers’ file names. In our experiment, noun types are determined by actually traversing the hierarchies, therefore they correspond to the top categories.




Figure 1: Part of type dependency graph for WordNet verbs

as Bayesian network learning (eg. Spirtes *et al.*, 1993). Those techniques use sophisticated methods that take into consideration of multiple antecedents/causations and so on, and build a complex and precise model with probabilities associated with edges. In our present work however, WordNet is compiled from human lexicographers' entries, thus the data has a fair amount of arbitrariness (i.e., noisy data). Therefore, we chose a simple technique which yields a simpler network, and used the result as a rough approximation of the type dependencies to be corrected manually at the next phase.

The advantage of this automatic analysis here is two fold: not only it discovers/reveals the semantic type associations with respect to the basic types selected from the previous step, it also helps the manual filtering to become more informed and consistent than by judging with mere intuition, since the result is based on the actual content of WordNet.

The type dependency graph is constructed in the following way. First, for all type-pairs extracted from the ambiguous words in WordNet, *mutual information* is computed to obtain the association by using the standard formula: for type t_1 , t_2 , a mutual information $I(t_1, t_2)$ is

$$I(t_1, t_2) = \lg \frac{f(t_1 \wedge t_2)}{\frac{f(t_1)}{N} \times \frac{f(t_2)}{N}}$$

where $f(t)$ is the number of occurrence of the type t , and N is the size of the data. The association between two types are considered prominent when the mutual information value was greater than some threshold (in our current implementation, it is 0).

At this point, type associations are undirected because mutual information is symmetric (i.e., commutative). Then, these associations are manually inspected to create a directed type dependency graph in the next phase. The manual filtering does two things: to filter out the spurious relations (i.e., false positives) and add back the missing ones (i.e., false negatives), and to determine the direction of the correct associations. Detected false positives are mostly homonyms (including metaphors) (eg.

WEA-EMO (weather and emotion) verb type pair for words such as the word ignite). False negatives are mostly the ones that we know exist, but were not significant according to the cooccurrence statistics (eg. ANI-FOOD in nouns). As a heuristic to detect the false negatives, we used the cross-categorical inheritance in the taxonomies in which category switches as the hierarchy is traversed up.

The direction of the associations are determined by *sense extension* described in section 2. In addition, we used the ontological generality of the basic types as another criteria. This is because a transitive inference through a ontologically general type may result in a relation where unrelated (specific) types are combined, particularly when the specific types are domain categories. For instance, the verb category CHA (change) is ontologically general, and may occur with specific types in entailment relation. But the transitive inference is done through this general type does not necessarily guarantee the systematicity between the associated specific types. In order to prevent this kind of implausible inference, we restricted the direction of a systematic relation to be from the specific type to the general type, if one of the member types is the generalization of the other. Note for some associations which involve equally general/specific types ontologically (such as COG (cognition) and COMM (communication)), the direction was considered bidirectional (unless sense extension strongly suggests the dependency). A part of the type dependency graph for WordNet verbs is shown in Figure 1.

3.3 Underspecified Semantic Classes

Underspecified semantic classes are automatically formed by partitioning the ambiguous senses of each word according to the type dependency graph.

Using the type dependency graph, all words in WordNet verb and noun categories are assigned one or more type *partitions*. A partition is an ordered set of basic types (abstracted from the fine-grained word senses in the first step) keyed by the primary type encompassing the secondary types. From a list of frequency-ordered senses of a WordNet word, a partition is created by taking one of the three most frequent types (listed as the first three senses in the WordNet entry) as the primary and collecting the secondary types from the remaining list according to the type dependency graph.³ Here, the secondary types are taken only from the nodes/types that are directly connected to the primary type. That is be-

³The reason we look at the first three senses is because primary types are not always listed as the most frequent sense in the WordNet sense lists (or in actual usage for that matter). We chose the first three senses because the average degree of polysemy is around 3 for WordNet (version 1.6) verbs and nouns.

Table 1: Example verbs in CONT classes

Verb Class	Verbs
CONT-CHA	<i>blend, crush, enclose, fasten, fold, puncture, tie, weld</i>
CONT-MOT	<i>beat, chop, fumble, jerk, kick, press, spread, whip</i>
CONT-POSS	<i>pluck, release, seize, sponge</i>
CONT-MOT-CHA	<i>dip, gather, mount, take_out</i>
CONT-MOT-POSS	<i>carry, cover, fling, toss</i>

cause we assumed if an indirect transitive dependency of $t1$ on $t3$ through $t2$ is strong enough, it will be captured as a direct dependency. This scheme also ensures the existence of a core concept in every partition (thus more implausible than transitive composition). This procedure is applied recursively if the sense list of a word was not covered by one partition (note in this case, the word is a homonym).

As an example, for the verb *write* whose sense list is (VCR COMM PCR CHA),⁴ the first 3 types VCR, COMM and PCR are picked in turn as the primary type to see whether a partition can be created that encompasses all other member types. In this case, a partition keyed by PCR can cover all member types (see the type dependency graph in Figure 1), thus a class VCR-COMM-PCR-CHA is created. The systematic relation of this class would be “a change or creation action which involves words (and resulted some object), performed for communication purpose (through the object)”.

For the verbs and nouns in WordNet 1.6, 136 underspecified verb classes and 325 underspecified noun classes are formed. Some verbs of the classes involving contact (CONT) are shown in Table 1. We can observe from the words assigned to each class that member types are indeed systematically related. For example, CONT-MOT class represents an action which involves physical contact resulting from motion (MOT). Words assigned to this class do seem to have motion flavor. On the other hand, CONT-POSS class represents a transfer of possession (POSS) which involves physical contact. Again, words in this class do seem to be used in a context in which possession of something is changed. For the more polysemous class CONT-MOT-POSS, words in this class, for instance *toss*, do seem to cover all three member types.

By using the underspecified classes, the degree of ambiguity in WordNet has substantially decreased. Table 2 shows the summary of our results (indicated by Und) compared to the original WordNet statistics. There, the advantage of our underspecified classes for reducing ambiguity seems very effective

⁴The original 9 senses listed in WordNet were compressed down to these 4 basic types.

Table 2: Average degree of ambiguity in WordNet

Category	All words		Polysemy only	
	WordNet	Und	WordNet	Und
verb	2.13	1.37	3.57	2.39
noun	1.23	1.06	2.73	2.21

for polysemous verbs (from 3.57 to 2.39, 33 % decrease). This is an encouraging result because many familiar (frequently used) verbs are polysemous in actual usage.

4 Application

To observe how the induced underspecified classes facilitates abductive inference in the contextual understanding of real-world texts, predicate-argument structures were extracted from the Brown corpus.⁵ Table 3 shows some examples of the extracted verb-object relations involving the verb class VCR (*verbal_creation*).

Abductive inference facilitated by underspecified classes is most significant when both the predicate and the argument are systematically polysemous. We call this a *multi-facet matching*.⁶ As an example, the verb *write* (VCR-COMM-PCR-CHA) takes an object noun *paper* (AFT-COMM) in a sentence in Brown corpus

In 1948, Afranio Do Amaral, the noted Brazilian herpetologist, wrote a technical paper on the giant snakes.

In this sentence, by matching the two systematically polysemous words *write* and *paper*, multiple interpretations are simultaneously possible. The most preferred reading, according to the hand-tagged corpus WNSEMCOR, would be the match between VCR of the verb (sense # 3 of *write* - to have something published, as shown in section 1) and COMM of the noun (sense # 2 of *paper* - an essay), giving rise the reading “to publish an essay”. However in this context, other readings are possible as well. For instance, the match between verb VCR and noun AFT (a printed media), which gives rise the reading “to have a written material printed for publishing”. Or another reading is possible from the match between verb COMM (sense # 2 of *write* - to communicate (thoughts) by writing) and noun AFT, which gives

⁵Predicate-argument structures (verb-object and subject-verb relations in this experiment) are extracted by syntactic pattern matching, similar to the cascaded finite-state processing used in FASTUS (Hobbs, *et al.*, 1997)). In the preliminary performance analysis, recall was around 50 % and precision was around 80 %.

⁶By taking the first sense for both predicate verb and argument noun, 78 % of the verb-object relations and 66 % of the subject-verb relations were systematically polysemous for at least one constituent.

Table 3: Examples of verb-object relations extracted from Brown corpus

Verb Class	Verb	Object Nouns
VCR	<i>pen</i>	note (COMM-ATT-POSS), dispatch (COMM-ACT-ATT)
VCR-MCR	<i>draft</i>	agreement (COMM-ATT-COG-REL-ACT), ordinance (COMM)
VCR-COMM	<i>write_out</i>	number (ATT-COMM), question (ACT-COMM-ATT)
	<i>dramatize</i>	comment (COMM-ACT), fact (COG-COMM-STA), scene (LOC)
VCR-COMM-PCR-CHA	<i>write</i>	article (AFT-COMM-ART-REL), book (AFT-COMM), description (COMM-ACT-COG), fiction (COMM), letter (COMM-ACT), paper (AFT-COMM), song (AFT-ACT-COMM)

rise the reading “to communicate through a printed media”. This reading implies the *purpose* and *entailment* of the *write* action (as COMM): a paper was written to communicate some thoughts, and those thoughts were very likely understood by the readers. Also from those readings, we can infer the *paper* is an *artifact*, that is, a physical object rather than an intangible mental object such as “idea” for instance. Those secondary readings can be used later in the discourse to make further inferences on the *write* action, and to resolve references to the *paper* either from the content of the paper (i.e., essay) or from the physical object itself (i.e., a printed artifact).

One interesting observation on multi-facet matching is the polysemous degrees of matched classes. Table 4 shows the predicate verbs of different systematically polysemous classes and the average polysemous degree of argument nouns observed in verb-object and subject-verb relations.⁷ The result indicates, as the verb becomes more polysemous, the polysemous degree of the argument stays about the same for both subject and object nouns. This suggests a complex multi-facet matching between verb and noun basic types, since the polysemous degree of nouns does not monotonically increase.

5 Discussion

The induction method described above should be considered as an initial attempt to automatically acquiring systematic polysemy from a broad-coverage lexical resource. The task is essentially to map our semantic/ontological knowledge about the systematicity of word meanings to some computational terms for a given lexical resource. In our present work, we mapped the systematicity to the cooccurrence of word senses. But the mapping only by computational/automatic means (mutual informa-

⁷The predicate-argument structures in this table represent the ones in which both verb and noun entries are found in WordNet. The total numbers of structures extracted from Brown corpus were 47287 for verb-object and 39266 for subj-verb. Discrepancies were mostly due to proper nouns and pronouns which WordNet does not encode.

tion) was not possible: manual filtering was further needed to enhance the mapping.

Also, there was a difficulty with type dependency graph. In the current scheme, systematicity among polysemous senses are represented by binary relations between a primary and a secondary sense in the graph. A partition, and eventually an under-specified class, is formed by taking all the secondary senses from the primary sense listed in each WordNet entry. The difficulty is that some combinations do not seem correct collectively. For example, a class PER-COG-CONT consists of two binary relations: PER-COG (to reason about what is perceived, eg. *detect*), and PER-CONT (to perceive through physical contact, eg. *hide*). Although each one correctly represents a systematic relation, PER-COG-CONT does not seem correct as a collection. In the WordNet entries, a verb *bury* is assigned to this class PER-COG-CONT. Here, CONT sense seems to select for a physical object (as in “They buried the stolen goods”), whereas COG sense (to dismiss from the mind) seems to select for a mental non-physical object. Therefore the construction of type partitions needs more careful considerations. Also the applicability of the induced classes must be evaluated in the further analysis.

6 Future Work

The work described in this paper is still preliminary. Our current induction method is semi-automatic, requiring some manual intervention. The first two steps, which selects basic types and creates type dependency graph, could be improved to further decrease the amount of manual effort, possibly to fully automated processes. The issues, then, will be how to detect the right level of abstraction and how to incorporate our linguistic knowledge as a prior domain knowledge in the induction algorithm for the given resource (WordNet).

Our next plan is to further analyze the result of the experiment and extract the selectional preferences, which will help disambiguate and *refine* the polysemous senses to a more restricted set of senses used in the context. However, as pointed out in (Resnik, 1997), strong selectional preferences may not be observed for broad-coverage texts, particu-

Table 4: Systematically polysemous verbs and average polysemous degree of argument nouns

Verb		Object		Subject	
Verb Class Poly Deg	# of Verbs	# of Nouns	Average Noun Class Poly Deg	# of Nouns	Average Noun Class Poly Deg
1	2729	9104	2.00	8969	1.71
2	714	5934	2.02	3884	1.65
3	169	2948	1.96	2402	1.72
4	34	1958	1.98	1640	1.71
5	1	279	1.95	87	1.37
Total	3647	20223		16982	

larly at the abstract level which our underspecified classes are defined.

Another important extension is to define a representation for each underspecified class that explicitly encodes how the senses relate to one another. Such information, which captures the implicit, complicated interactions between different aspects of an action which may involve implied objects, can be encoded in a structured lexical representation that is along the same line of some recent research in lexical semantics (eg. Pustejovsky, 1995; Verspoor, 1997) and knowledge representation. Then, it will be interesting to see such representation defined at the abstract polysemous class level can be combined with micro (word sense) level representation (eg. (Harabagiu and Moldovan, 1997)).

Acknowledgments

The author would like to thank Paul Buitelaar for helpful discussions, insights and encouragement.

References

- Apresjan, J. (1973). Regular Polysemy. *Linguistics*, (142).
- Buitelaar, P. (1997). A Lexicon for Underspecified Semantic Tagging, In *Proceedings of the ACL SIGLEX Workshop on Tagging Text with Lexical Semantics*. pp. 25-33.
- Buitelaar, P. (1998). CORELEX: Systematic Polysemy and Underspecification. Ph.D. dissertation, Department of Computer Science, Brandeis University.
- Copestake, A. and Briscoe, T. (1992). Lexical Operations in a Unification-based Framework, In *Lexical Semantics and Knowledge Representation*, J. Pustejovsky (eds.), pp. 101-119, Springer-Verlag.
- Copestake, A. and Briscoe, T. (1995). Semi-productive Polysemy and Sense Extension. *Journal of Semantics*, 12.
- Fellbaum, C. (1990). English Verbs as a Semantic Net. *International Journal of Lexicography*, 3 (4), pp. 278-301.
- Harabagiu, S. and Moldovan, D. (1997). TextNet - A Text-based Intelligent System. *Natural Language Engineering*, 3.
- Hobbs, J., Appelt, D. Bear, J., Israel, D., Kamayama, M., Stickel, M. and Tyson, M. (1997). FASTUS: A Cascaded Finite-state Transducer for Extracting Information from Natural-language Text, In *Finite-state Language Processing*, E. Roche and Y. Schabes (eds.), pp. 383-406, The MIT Press.
- Kilgarriff, A. (1997). Foreground and Background Lexicons and Word Sense Disambiguation for Information Extraction, In *Proceedings of the International Workshop on Lexically Driven Information Extraction*.
- Miller, G. (eds.) (1990). WORDNET. An Online Lexical Database. *International Journal of Lexicography*, 3 (4).
- Ostler, N. and Atkins, B. (1992). Predictable Meaning Shift: Some Linguistic Properties of Lexical Implication Rules, In *Lexical Semantics and Knowledge Representation*, J. Pustejovsky (eds.), pp. 87-100, Springer-Verlag.
- Procter, P. (1978). *Longman dictionary of Contemporary English*, Longman Group.
- Pustejovsky, J. (1995). *The Generative Lexicon*, The MIT Press.
- Resnik, P. (1997). Selectional Preference and Sense Disambiguation, In *Proceedings of the ACL SIGLEX Workshop on Tagging Text with Lexical Semantics*. pp. 52-57.
- Spirites, P., Glymour, C. and Scheines, R. (1993). *Causation, Prediction and Search*, Springer-Verlag.
- van Deemter, K. and Peters, S. (1996). *Semantic Ambiguity and Underspecification*, CSLI Lecture Notes 55, Cambridge University Press.
- Verspoor, C. (1997). Contextually-dependent Lexical Semantics. Ph.D. dissertation, University of Edinburgh.