

Towards Exemplar-based Polysemy

Mohsen Rais-Ghasem (mohsen@cyberus.ca)
Jean-Pierre Corriveau (jeanpier@scs.carleton.ca)

School of Computer Science, Carleton University
Ottawa, ON K1S 5B6 Canada

Abstract

In this paper we criticize existing computational models of lexicon for assuming that for every word there is a fixed number of word sense that must be searched for the proper meaning of that word in a context. We reject this sense enumerative view and argue for a different model of lexicon in which the effects of context are not limited to selecting a word sense, and selected senses can be contextually modulated. We also explain how patterns of contextual effects could evolve in an exemplar-based fashion. A prototype implementation of this model is also discussed.

Introduction

This paper proposes a computational model of lexicon. What distinguishes the proposed model from other such models is the treatment of polysemy. Most existing models of lexicon presuppose a list of possible word senses for each word, and a selection process that searches such a list for the proper meaning of that word in a context. We question this sense enumerative approach and argue for an alternative model in which the ultimate semantic characteristics of a word not only depend on the meaning of the word itself but on the interactions that the word has with other words in its contexts. We will also show how the model can gradually extract patterns of contextual effects from a number of exemplars.

We will first review computational models of lexicon and argue that they essentially present an enumeration of senses in one form or another. We will then question the validity and viability of such an approach. Following this section, we will present our model and explain how word senses in our model are determined. Finally, a prototype implementation of the proposed model and results of some sample runs will be presented.

Overview of Computational Lexicons

Computational models of lexicon can be divided into three categories: symbolic lexicons, structured connectionist lexicons, and distributed connectionist lexicons. In symbolic lexicons, word senses are generally represented by means of *ad hoc* and complicated structures. Typically, such structures are predefined, i.e. hand-coded, and static. Lexicons of most early NLP models (e.g. BORIS, Dyer 1983) fall under this category.

In structured connectionist lexicons (e.g.; Bookman, 1994; Lange and Dyer, 1989; Waltz and Pollack, 1985) each word is represented by a node in a structured network. Word nodes become activated when their corresponding words are presented at the input. Each word node is connected to some sense nodes, each representing one of the word's senses. For example, in the network described in (Waltz and Pollack, 1985) the word *star* is connected to three senses: *movie star*, *celestial body*, and *geometric figure*.

Distributed connectionist networks are the basis of the last category of computational lexicons. Unlike structured connectionist lexicons where every concept is assigned to a node, here a single set of nodes is used to represent all word senses. Examples of such lexicons are reported in McClelland and Kawamoto (1986), Miikkulainen (1993), and Veronis and Ide (1995). For example, in McClelland and Kawamoto (1986), each word is associated with one or more word senses. Word senses are encoded along a number of semantic microfeatures. Examples of these microfeatures are softness (soft/hard), gender (male/female/neuter), volume (small/medium/large), and form (compact/1-D/2-D/3-D).

Despite the difference in their adopted concept representation approach, all three categories of computational lexicons are committed to the same principle of enumerating word senses¹. This sense enumerative view postulates that, upon hearing a word, listeners access a mental dictionary containing an exhaustive list of potential senses for that word, from which, the proper sense for the word is selected.

On Sense Enumeration

Sense enumeration seems very intuitive. After all, it is how all dictionaries are organized. The question seems to be about the validity of accepting this familiar assumption as a meaning theory. Ruhl remarks:

I claim that current linguistic theories accept too uncritically the conclusion of dictionaries that words in general have multiple meanings (Ruhl, 1989:vii).

Kilgarriff (1997) contends that standard dictionaries specify the range of meanings of a word in a list merely in response

¹ While in most symbolic and structured connectionist model this commitment is made at the architectural level, in distributed connectionist model it is mostly a matter of representation.

to “constraints imposed by tradition, the printed page, compactness, a single, simple method of access, and resolving disputes about what a word does and does not mean.” He criticizes word sense disambiguation (WSD) approaches for committing themselves to such lists:

Much WSD word proceeds on the basis of there being a computationally relevant, or useful, or interesting, set of word senses in the language, approximating to those stated in a dictionary... Meanwhile, the theoreticians provide various kinds of reasons to believe there is no such set of senses (Kilgarriff, 1998:95).

Moreover, it is not always possible to enumerate all potential senses of a word. The use of language, even non-figuratively, by its very nature is creative. That is, “words can take on an infinite number of meanings in novel contexts” (Pustejovsky, 1995: 42). No matter how comprehensive a list can be, it is generally possible to find a new sense for some words. *Eponymous* expressions, expressions built around references to people, provide a good example. Clark and Gerrig (1983) showed how understanding eponymous expressions such as *Please do a Napoleon for the camera* requires sense generation rather than sense selection.

Another problem facing sense enumeration is *semantic flexibility*. Many psycholinguistic findings suggest that context highlights (or obscures) certain properties of a single concept as it appears in different contexts. For example, in an early experiment, Barclay et al. (1974) demonstrated how the interpretations of familiar, unambiguous words vary with context. For instance, they argued that the choice of attributes for *piano* is affected by the verb selection in *The man (lifted) (tuned) (smashed) (sat on) (photographed) the piano*. They then provided evidence that the prior acquisition of a sentence like *The man lifted the piano* (vs. *The man tuned the piano*) influences the effectiveness of cues like “something heavy” (vs. “something with a nice sound”) in recall. They concluded that context can affect the encoding of concepts in memory. Similar results have been reported in Anderson et al. (1976), Barsalou (1982), Greenspan (1986), and Witney et al. (1985).

Having subscribed to the sense enumeration principle, the effect of context in existing models of lexicon is limited to selecting a sense among a closed number of alternatives. Any changes in the characteristics of a selected sense either has to come in the form of a new sense or is ignored.

Dealing with contextual effects is particularly problematic for symbolic and structured connectionist lexicons. This is due to the fact that commitment to sense enumeration is made at the structural level, making word senses discrete entities. To the contrary, word senses in distributed connectionist models are continuous. Thus, such models are potentially capable of dealing with contextual effects. In fact McClelland and Kawamoto (1986) reported an unintended yet interesting result. They had presented their model with *The ball broke the vase*. Although throughout the training phase *ball* was always associated with the microfeature *soft*, in the output it was associated with the microfeature *hard*. They attributed this result to the fact that *breakers* in their experiment were all *hard* and the model had shaded the

meaning of *ball* accordingly. Kawamoto later remarked:

[T]he flexibility of a distributed representation allows a natural account of polysemy and homonymy, and provides a mechanism for new senses to be learned or generated on-line (1993: 510).

However, distributed connectionist approach also has some disadvantages. First of all, such lexicons presuppose a set of universal and fixed microfeatures and demand every sense to be characterized in terms of such a set *in advance*. But what is even more important is the difficulty to separate patterns of contextual effects from the representation of a word sense. For instance, consider *breakers* in McClelland and Kawamoto (1986). It is impossible to isolate *breakers* (the category) from *breakers* (the instances).

We believe this separation is useful. Firstly, such patterns can be thought of as *ad hoc* categories: categories built by people to achieve goals (Barsalou, 1983). For instance, *breakers* can be instrumental in achieving the goal of “breaking a window”. Secondly, from a learning point of view, such patterns can be very useful. Rais-Ghasem (1998) has shown how a concept can evolve (i.e. acquire new properties) from such patterns as it becomes associated with in different contexts. And finally, Rais-Ghasem (*Ibid.*) has employed such patterns to implement a metaphor understanding theory, in which metaphors are interpreted as class inclusion assertions (Glucksberg & Keysar, 1990).

A Lexicon for Sense Modulation

In this section we will discuss our proposed model of lexicon. We begin by introducing our two-tiered word senses. Later in this section, we explain how this two-tiered structure allows us to account for contextual effects in an exemplar-based fashion. The structure and overall behavior of the model are described in the last two subsections.

Two-Tiered Word Senses

Word senses in the proposed model consist of two components: **sense-concept** and **sense-view**. Given a word sense, its sense-concept will determine the concept it represents, whereas, the sense-view describes how this sense-concept is to be viewed in its surrounding context. For example, the word sense for *book* in *The book broke the window* will consist of the sense-concept BOOK, representing the concept book, and a sense-view which portrays *book* as an instrument of breaking.

This separation parallels the two different roles that Franks (1995) proposed for concepts when he distinguished between their *representational* and *classificatory* functions. While the former is used to discern instances of one concept from instances of others, the latter specifies how an instance of a concept should be classified. For example, he argues that, depending on context, *fake gun* could be classified as a *gun*, a *toy*, a *replica*, and a *model*.

We contend that this two-tiered structure allows us to account for various contextual effects. In general, Cruse (1986) specifies two ways in which context affects the semantic contribution of a word: *sense selection* and *sense*

modulation. Sense selection happens in cases of lexical ambiguity where one sense is chosen among a number of distinct senses. Sense selection is followed by sense modulation in which the semantic characteristics of selected senses are modulated or become specified according to their surrounding contexts.

It is our intention to show that unlike conventional lexicons, which do not go beyond sense selection, using the proposed two-tiered structure for word senses we can account for both processes. Below we will show that sense selection and sense modulation, respectively, lead to the selection of a sense-concept and a sense-view. As an example, consider the word *book*. Being unambiguous, the sense-concept BOOK can be easily selected. However, BOOK must be modulated according to its surrounding context. For example, in a context such as *The book broke the window* it will be associated with a sense-view that portrays it as breaker. In other contexts such as *Many books were burnt in the fire* and *I read the book* the same concept BOOK would be associated with different sense-views namely a flammable object and text.

Development of Sense-Views

Concepts associated with a word, to a large extent are conventionalized. Traditionally, concepts are represented by means of a number of weighted properties (for a review see Barsalou & Hale, 1993). Such properties explain similarities and facilitate comparisons. While representation of concepts and word-concept mapping are pre-defined in the proposed model, sense-views are developed incrementally in an exemplar-based fashion.

This is achieved by defining the *alike* relationship. Essentially, this relationship states that two sense-concepts become alike if they appear in a similar context and they share the same thematic role. For example, *book* in *The book broke the window* and *bat* in *The man smashed the windshield with a bat* are alike. That is to say that, *book* in this context is best classified with *bat* and not with *book* in a context such as *I read the book*, despite the fact that representationally it should be the other way around.

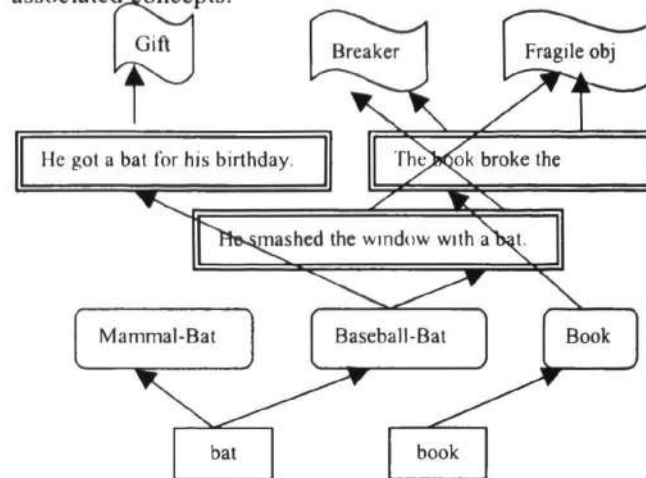
Now considering the two-tiered word senses, we can say that sense-concepts associated with a sense-view are alike. Thus, a sense-view can be considered as a generalization of its associated sense-concepts. Therefore, as new sense-concepts become associated with a sense-view, its characteristics will evolve to become a better representative of all its associated sense-views (see Sample Runs below for examples).

A Lexicon with Exemplars

The proposed model is structured in four levels. Words appear at the bottom level. Concepts associated with words appear at the second level. Unlike other models, the proposed model also maintains examples of how each concept is typically used. Such usage examples, or *exemplars*, appear at the third level. Each exemplar consists of a number of sense-concepts, each representing the occurrence of a concept in the context given by the

exemplar. Sense-views constitute the fourth level. Each sense-concept is connected to a sense-view.

The following picture illustrates an example of two words, *bat* and *book* (represented as rectangles) and their associated concepts (represented as rounded rectangles). Concepts are further specified by exemplars (represented as double-lined rectangles) each representing a use-case for its associated concepts.



For the sake of simplicity, sense-concepts constituting each exemplar are not displayed. An example of sense-concept is the occurrence of BASEBALL-BAT in *He got a bat for his birthday*. Each sense-concept is connected to a sense-view. Three sense-views are shown: breaker, gift and fragile object. Note that both sense-views and sense-concepts are in fact represented by a set of properties, and that the names used in this figure are only labels used for the sake of clarity and brevity.

Finally, here is the output generated by the implemented prototype (see Implementation) for *book* in *The book broke the window*.

```
SENSE Generated for Input Word Book
[Sense-Concept(s)] BOOK
[Sense View]
Thematic Role: Instrument   No. of Exemplars: 4
Marker(s): with-pp, p-subj.
STATE-OF-MATTER(0.73305)---->: SOLID,
MADE-OF(0.82805)---->: MATERIAL,
```

As shown, the sense-view breaker consists of two properties with strong presence (enclosed number represents weight, between 0 and 1, for each property). Thus, this output should be interpreted to indicate that, in the given context, BOOK is to be viewed as a solid object made of material. In other words, other properties of BOOK such as author or title, as well as the fact that book is made of paper, ink etc., are all irrelevant in this context.

From Words to Senses

Criticizing Lakoff for concluding that *open* is polysemous because of expressions such as *open the door* and *open the present*, Ruhl writes:

Admittedly, the phrase *open the door* and *open the present* evoke quite different images, but why is the

difference attributed to *open*, which does not differ, and not also *the?* Why isn't the difference located solely where there is difference: in the door-present distinction and in the knowledge that people have about these two activities (1989:x).

This notion that the meaning of a word is being dynamically made up, at least in part, of meanings of other words it occurs with in a context constitutes the basis of the process that maps input words to word senses in the proposed model².

The process begins by presenting an input context, made up of input words to the model. The goal is to find pairs of sense-concept/sense-view for input words. Details of this process are beyond the scope of this paper (see Rais-Ghasem, 1998). What follows is a brief overview of the process.

Word nodes corresponding to the input words are activated. Activated word nodes, in turn, activate their related concepts, which in turn activate their related exemplars. The set of activated exemplars represents the model's knowledge, up to that point, of various ways that input words can interact with other words.

Having located all such possible interactions for every input word, the next step is to find a set of senses for input words that can **interact** with each other. A number of word senses can interact with each other only if they either appear in one of the exemplars, or they can be used in one.

The decision on whether input words can be interpreted based on one of the exemplars depends on the possibility of classifying the input words along that exemplar. In other words, the question is whether a word sense can be found for each input word so that it could be associated with the same sense-view that its corresponding sense-concept in the exemplar is associated with.

For example, assume that *The book broke the window* is presented to the model and two exemplars are activated for this input: *John broke the law* and *The man smashed the windshield with a bat*. The first exemplar is rejected because none of the concepts associated with the word *window* could be classified as law and regulation (the sense-view related to sense-concept LAW). Because senses can be found for input words that are compatible with the sense-views associated with the second exemplar, the input sentence can be interpreted after it. For example, *book* (in fact its related concept BOOK) can be classified as a breaker. Note that finding such set of senses implies finding a pair of sense-concept/sense-view for each input word.

Finally, the examination of activated exemplars can be carried out in parallel. In fact, in the next section, we discuss a prototype implementation in which each activated exemplar measures its adaptability to the input.

Implementation

This section reports on a prototype implementation of the proposed model for sentential contexts. For this

² Also see *semantic traits* in Cruse, 1986.

implementation, we used syntactic case markers (Delisle et al., 1993) as indications of thematic roles that noun phrases play in a context. Examples of such case markers are p-subj (positional subject) and with-pp (prepositional phrase with).

Architecture

The implemented prototype is a hybrid system. Architecturally, the system is grounded in two marker passing networks. The bottom network, the ontology network, serves as the system's knowledge base to define concepts in the second network. The network is arranged based on the Mikrokosmos ontology (Mahesh & Nirenburg, 1995). The second network, the lexicon network, consists of four layers of nodes: words, concepts, exemplars comprising sense-concepts, and sense-views (as discussed earlier).

However, despite the fact that it is built on top of two structured networks, the behavior of the system is not completely accomplished through a spreading activation process. Instead, its functionality to a large extent arises from the execution of a number of processing elements called *agents*. Attached to network nodes, agents are autonomous processing elements capable of performing a sequence of instructions. Once a node is activated, its attached agent "fires up" and individually and concurrently starts carrying out its instructions. Details of types and responsibilities of agents employed in this implementation could be found in Rais-Ghasem (1998).

Because of its hybrid nature, the implemented system supports parallelism at the knowledge level, the ability to apply symbolic rules and controls, and taking advantage of statistical similarities in developing sense-views.

Sample Runs

Results of some of the test cases conducted on the implemented system are presented in this subsection.

Sense Modulation: This test demonstrates the system's capability to modulate semantic features of one single sense-concept in different contexts. Here is the first input context:

The musician moved the piano.

The output sense generated for piano is displayed below. Note how only those properties of PIANO that portray it as a heavy physical object are highlighted by the sense-view.

```
SENSE Generated for Input Word Piano
[Sense-Concept(s)] PIANO
[Sense View]
Thematic Role: Object - No. of Exemplars: 3   Marker(s):
p-obj,
WEIGHT(0.666667)---->:heavy,
IsKindOf-DEVICE(0.513)---->:
STATE-OF-MATTER(0.756)---->:SOLID,
MADE-OF(0.513)---->:PLASTIC,METAL,
IsKindOf-ARTIFACT(0.7047)---->:
COLOR(0.7047)---->:
AGE(0.7047)---->:
OPERATED-BY(0.7047)---->:HUMAN,
```

It is also interesting to look at the selected sense-view for *musician* in which all properties specific to musician are suppressed, since they are irrelevant in the present context.

SENSE Generated for Input Word Musician
 [Sense-Concept(s)] MUSICIAN
 [Sense View]
 Thematic Role: Agent - No. of Exemplars: 3 - Marker(s):
 p-subj,
 GENDER(0.885367)---->:MALE,
 IsKindOf-HUMAN(0.8187)---->:
 IsKindOf-PRIMATE(0.54)---->:

Here is the second input context:

The musician played the piano.

In contrast, here the properties portraying *piano* as a musical instrument and musician are relevant.

SENSE Generated for Input Word Piano
 [Sense-Concept(s)] PIANO
 [Sense View]
 Thematic Role: Object - No. of Exemplars: 2 Marker(s):
 p-obj,
IsKindOf-MUSICAL-INSTRUMENT(0.81)---->:
WORK-EQUIPMENT-OF(0.81)---->:MUSICIAN,
 IsKindOf-ARTIFACT(0.729)---->:
 STATE-OF-MATTER(0.729)---->:SOLID,
 COLOR(0.729)---->:
 AGE(0.729)---->:
 OPERATED-BY(0.729)---->:HUMAN,
 IsKindOf-INANIMATE(0.6561)---->:

Here is the last input context and generated sense:

The man broke the piano.

SENSE Generated for Input Word Piano
 [Sense-Concept(s)] PIANO
 [Sense View]
 Thematic Role: Object - No. of Exemplars: 4 Marker(s):
 p-obj,
 IsKindOf-ARTIFACT(0.54675)---->:
 STATE-OF-MATTER(0.567)---->:SOLID,
 COLOR(0.54675)---->:
 AGE(0.54675)---->:
 OPERATED-BY(0.54675)---->:HUMAN,
 MADE-OF(0.6775)---->:PLASTIC, GLASS,

Sense-View Development: This test demonstrates gradual development of sense-views. The destination sense-view is initially exemplified by only one exemplar:

Mary went to the office.

[Sense View]
 Thematic Role: Destination - Marker(s): to-pp,
 IsKindOf-BUILDING(0.9)---->:
 IsKindOf-PLACE(0.81)---->:
 IsKindOf-PHYSICAL-OBJECT(0.729)---->:
 MADE-OF(0.729)---->:MATERIAL,
 WEIGHT(0.729)---->:
 SIZE(0.729)---->:
 IsKindOf-OBJECT(0.6561)---->:
 IsKindOf-BUILDING-ARTIFACT(0.81)---->:
 IsKindOf-ARTIFACT(0.729)---->:
 STATE-OF-MATTER(0.729)---->:SOLID,
 COLOR(0.729)---->:
 AGE(0.729)---->:
 OPERATED-BY(0.729)---->:HUMAN,
 IsKindOf-INANIMATE(0.6561)---->:

Notice both IsKindOf-Building and IsKindOf-Place are relatively central to office and therefore to this sense-view.

The set properties shrinks rapidly after processing the following exemplar:

The student went to the stadium.

SENSE Generated for Input Word Stadium
 [Sense-Concept(s)] STADIUM
 [Sense View]
 Thematic Role: Destination - Marker(s): to-pp,
 IsKindOf-BUILDING(0.8145)---->:
 IsKindOf-PLACE(0.73305)---->:
 IsKindOf-BUILDING-ARTIFACT(0.73305)---->:

This trend continues with another input:

John went to the park.

SENSE Generated for Input Word Park
 [Sense-Concept(s)] PARK
 [Sense View]
 Thematic Role: Destination - Marker(s): to-pp,
 IsKindOf-BUILDING(0.543)---->:
 IsKindOf-PLACE(0.7074)---->:

Here, unlike previous case, the property IsKindOf-Place is more prominent than IsKindOf-Building. This is because *park* is not a building, nonetheless, its effect is not enough to completely eliminate IsKindOf-Building from the sense-view.

Multiple Word Senses: There are cases in which context does not favor any of the alternative readings of a word, and therefore the ambiguity must be maintained in the output. This test demonstrates the system's ability to handle such cases. In this example, both readings of *bank* are compatible, to some degree, with the destination sense-view.

John went to the bank.

Here is the output word sense for *bank*, with two sense-concepts, both linked to the same sense-view.

SENSE Generated for Input Word Bank
 [Sense-Concept(s)] RIVER-BANK, BANK-BRANCH
 [Sense View]
 Thematic Role: Destination - Marker(s): to-pp,
 IsKindOf-BUILDING(0.51585)---->:
 IsKindOf-PLACE(0.5967)---->:

Unknown Words: Here is an example of how sense-views can be used to establish some properties about unknown words. Here is the input:

Mary went to the palladium.

The word *palladium* is not defined in the lexicon. Nevertheless, the system associates it with the proper sense-view. Through this sense-view, some initial properties for *palladium* can be inferred.

SENSE Generated for Input Word Palladium
 [Sense-Concept(s)] *** unknown ***
 [Sense View]
 Thematic Role: Destination - Marker(s): to-pp,
IsKindOf-BUILDING(0.51585)---->:
IsKindOf-PLACE(0.5967)---->:

Instantiation of General Terms: This test is inspired by the experiment reported by Anderson et al. (1976). They found that *shark* was a better cue than *fish* for subjects in

remembering a sentence like the following:

The fish attacked the man.

They concluded that *fish* was instantiated to, and encoded accordingly as, *shark* in their subjects' memory. Notice how in the output, *fish* is associated with properties specific to shark (*aggressiveness* and *black color*).

SENSE Generated for Input Word Fish
[Sense-Concept(s)] FISH
[Sense View]
Thematic Role: Agent - Marker(s): p-subj,
COLOR(1)---->:BLACK,
AGGRESSIVE(1)---->:
IsKindOf-FISH(0.9)---->:
IsKindOf-VERTEBRATE(0.81)---->:
IsKindOf-ANIMAL(0.729)---->:
GENDER(0.729)---->:
IsKindOf-ANIMATE(0.6561)---->:

Conclusion

In this paper we discussed a lexicon model in which the role of context is not limited to sense selection. Selected senses are also modulated according to their surrounding context. We also described how patterns of contextual effects could be learned by the model. A prototype implementation of the model was also discussed.

Acknowledgments

Support from the Natural Science and Engineering Research Council of Canada is gratefully acknowledged.

References

- Anderson, R., Pichert, J., Goetz, E., Schallert, D., Stevens, K., & Trollip, S. (1976) Instantiation of general terms. *Journal of Verbal Learning and Verbal Behavior* 15:667-679.
- Barclay, J., Bransford, J., Franks, J., McCarrell, N. & Nitsch, K. (1974) Comprehension and semantic Flexibility. *Journal of Verbal Learning and Verbal Behavior* 13:471-481.
- Barsalou L.W.(1982). Context-independent and context-dependent information in concepts. *Memory and Cognition* 10(11):82-93
- Barsalou L.W.(1983). Ad hoc categories. *Memory and Cognition* 11(3):211-227.
- Barsalou, L.W. & Hale C.R. (1993). Components of conceptual representation: from feature lists to recursive frames. In Mechelen I.V. et al, A. (Eds.), *Categories and concepts* Academic Press.
- Clark, H. & Gerrig, R. (1983). Understanding old words with new meanings. *Journal of Verbal Learning and Verbal Behavior* 22:591-608.
- Cruse, D. (1986). *Lexical Semantics*, Cambridge University Press.
- Delisle S., Copeck, T., Szpakowicz, S. & Barker, K. (1993). Pattern matching for case analysis: A computational definition of closeness. ICCL, 310-315.
- Dyer, M. (1983). *In-depth Understanding: A computer model of integrated processing for narrative comprehension*. Cambridge, MA: MIT Press.
- Franks, B. (1995). Sense Generation: A "Quasi-Classical" Approach to Concepts and Concept Combination. *Cognitive Science* 19:441-505.
- Glucksberg, S. & Keysar, B. (1990). Understanding metaphorical comparisons: Beyond literal similarity. *Psychological Review* 97(1): 3-18.
- Greenspan, S. (1986). Semantic flexibility and referential specificity of concrete nouns. *Journal of Memory and Language* 25:539-557.
- Kawamoto, A.H. 1993. Nonlinear dynamics in the resolution of lexical ambiguity: a parallel distributed processing account. *Journal of Memory and Language* 32:474-516.
- Kilgarriff, A. 1997. What is word sense disambiguation good for?. Proceedings of Natural Language Processing in the Pacific Rim(NLPRS-97) Phuket, Thailand.
- Kilgarriff, A. 1998. "I do not believe in word senses". *Computers and Humanities* 31(2): 91-113.
- Lakoff, G. (1987). *Women, Fire and Dangerous Things*. The University of Chicago Press.
- Lange, T. & Dyer, M. (1989). Frame selection in a connectionist model of high-level inferencing. *Proceedings of the 11th Conference of the Cognitive Science Society*.
- Mahesh K. & Nirenburg, S.(1995). A situated ontology for practical NLP. *Proceedings of the Workshop on Basic Ontological Issues in Knowledge Sharing*. IJCAI 95. Montreal, Canada.
- McClelland, J. & Kawamoto, A. (1986). Mechanisms of sentence processing. In McClelland J. and Rumelhart, D. (Eds.) *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. Vol. 2. MIT press.
- Miikkulainen, R. 1993. *Subsymbolic natural language processing: An integrated model of scripts, lexicon, and memory*. Cambridge, MA: MIT Press.
- Pustejovsky, J. (1995). *The generative lexicon*. MA, MIT Press.
- Rais-Ghasem,, M. (1998) *An exemplar-based account of contextual effects* (Ph.D. Thesis) Ottawa, ON: Carleton University, School of Comp. Sc.
- Ruhl, C. (1989) *On monosemy: A study in linguistic semantics* Albany, N.Y: State University of New York Press.
- Veronis, J. & Ide, N. (1995). Large neural networks for the resolution of lexical ambiguity. In Saint-Dizier, P. & Viegas, E. (Eds.) *Computational Lexical Semantics* NY, Cambridge University Press.
- Waltz, D. & Pollack, J. (1985). Massively parallel parsing: A strongly interactive model of natural language interpretation. *Cognitive Science* 9:51-74.
- Witney, P., McKay, T., & Kellas, G. (1985). Semantic activation of noun concepts in context. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 11:126-135.