Clustering in the Discovery of Semantic Frames

Rebecca Green
Hornbake Building (So. Wing), Room 4105
College of Information Studies
University of Maryland
rgreen@umd.edu

1. Introduction

This paper describes the use of clustering at three stages within a larger research effort to identify semantic frames used in English automatically.[1] The first of two tasks within this effort has been the identification of sets of semantically related verb senses that invoke a common semantic frame. Within this task, clustering has been used both to build sets of verb senses with the potential of invoking a common semantic frame and then to merge sets with a high degree of overlap.

The paper is organized as follows: Section 2 introduces frame semantics. Section 3 outlines the methodology used to identify sets of semantically related verb senses that invoke a common semantic frame, while section 4 presents the specific clustering algorithm used within that process. Section 5 discusses the use of this clustering algorithm for the identification of semantically related verbs in two machine-readable lexical resources: the machine-readable version of the *Longman Dictionary of Contemporary English* (LDOCE, 1978 edition) and WordNet, an online lexical database (http://www.cogsci.princeton.edu/~wn; version 1.7.1 has been used for the work reported here). Section 6 presents the use of clustering to merge overlapping sets of verb senses formed in previous steps. Section 7 discusses the results of these clusterings, paying particular attention to the effect of LDOCE's restricted defining vocabulary on the clustering process.

2. Frame Semantics

Frame semantics is a theory of language understanding, situated within a world in which some types of experiences—conditions, relationships, events—recur frequently and are structured. For example, one semantic frame corresponds to the condition of being physically tired; another semantic frame corresponds to the equivalence relationship; a third semantic frame corresponds to the event of playing a baseball game. Fillmore (1982) describes a semantic frame as "any system of concepts related in such a way that to understand any one of them you have to understand the whole structure in which it fits" (p. 111); he further describes a semantic frame as "a system of categories structured in accordance with some motivating context" (p. 119). For example, our understanding of physical weariness would be incomplete without including an understanding of effort/work, on the one hand, and of rest/sleep, on the other hand. The body's natural sleep cycle and the physical necessity of expending effort to procure that which is required to sustain and enhance human life provide the motivation for such a semantic frame. Likewise, we cannot fully understand a baseball game without understanding things like teams, innings, pitching, fly balls, strikes, fouls, ground rule doubles, and so forth. The motivating context for this frame is the entertainment value of playing or watching baseball.

---

[1]The hypotheses driving this research are: (1) that the use of semantic frames in knowledge-intensive tasks can improve performance (this is investigated with regard to text segmentation) and (2) that such frames can be induced from data in machine-readable lexical resources.

The recurrence of prototypical experiences and the cultural or institutional meaning they have motivates the introduction of linguistic expressions to refer to the whole of these experiences, as well as to their salient parts or aspects. Such words or phrases are said to "invoke" the frame and inherit its organization; in turn the frame structures the meaning of those words and phrases (Fillmore, 1982, p. 117). While any given word will typically highlight (or "profile") particular aspects of the frame, the whole of the frame structure is conveyed by each word or phrase that invokes and inherits it. For example, the phrase "full count" invokes the entire baseball frame, but profiles the relationship between pitcher and batter in a particular at-bat.

Frame semantics has great import for any task involving meaning. Within information retrieval, for example, the existence and properties of semantic frames mean that a specific type of experience may be discussed using any of the words and phrases that invoke the corresponding frame(s). Bag-of-words approaches to document retrieval will almost always fail to retrieve at least some of the most relevant documents available because of a failure to account for the full set of words and phrases used to invoke a semantic frame. Frame-semantic-aware approaches will be needed as part of the retrieval arsenal if we are to achieve optimal performance at this area. Likewise frame semantic awareness has a role to play in such other knowledge-intensive tasks as text summarization, machine translation, information extraction, topic identification, and text segmentation.

In the case of events, for which the richest semantic frames exist, the internal structure of the frame typically includes some number of roles (that is, participants in the event, identified by the functions they play in the event) and may also include attributes of the event and subevents of the event. From a linguistic point of view, these roles/participants are often expressed as the nominal arguments of frame-invoking verbs. For example, in the sentence "The catcher tagged the runner out at the plate", the baseball frame is invoked by *catcher, tagged, runner, out,* and *plate*; *tagged* profiles the particular event that occurred, *out* profiles the resulting condition of the event, *plate* highlights where the event took place, and *catcher* and *runner* name the participants.

One approach to achieving frame semantic awareness starts by identifying semantic frames both extensionally and intensionally. Extensional identification consists of enumerating all words (or more specifically, all word senses) that invoke the frame; intensional identification consists of detailing the internal structure of the frame. The work reported here is part of a research effort to accomplish the rudiments of such an identification process. The work has adopted several limitations: Only frame-invoking verb senses have been identified in the extensional task; only roles expressible by nouns have been identified in the intensional task. Thus, for example, the extensional task should find that specific senses of such verbs as *buy, purchase, sell, cost, pay, spend, expend, charge,* and *price* all invoke a Commercial Transaction frame. The intensional task should find that the internal structure of the frame includes a Buyer, a Seller, Merchandise, and Money (Fillmore, 1982, p. 116). Senses of *risk, venture, hazard, chance, gamble,* and *endanger* likewise invoke a Risk frame. The extended internal structure of this frame includes a degree of Chance that some Harm will befall a Victim or a Valued Object of the Victim because of a (Risky) Situation brought about by a Deed performed by an Actor for an (Intended) Gain or Purpose, on behalf of a Beneficiary, or impelled by a Motivation (Fillmore & Atkins, 1992, pp. 80-84).

3. Identifying Semantically Related Verb Senses: Methodology

The overall approach being undertaken in the extensional task is to gather weighted evidence from a variety of sources on sets of verb senses that are likely to invoke a common semantic frame. Pairs of verb senses supported by a certain number of data sources and/or whose total weight exceeds a specific threshold are carried forward and combined into larger groupings.

Among the major kinds of evidence used to posit semantic relatedness between verb senses are (1) shared subject codes in LDOCE, (2) explicit semantic links between entries in lexical resources (e.g., the hyperonymy/hyponymy, synonymy, antonymy, entailment, and causative relationships within WordNet and the COMPARE and OPPOSITE links within LDOCE), and (3) co-occurrence of, or relatedness between, the words used in their definitions. In this latter case, clustering has been used to determine which verb senses are most similar. The process of clustering verb senses based on words in their definitions has been applied to both LDOCE and WordNet. Data on pairs of verb senses related through these various sources and techniques are analyzed to generate groupings of verb senses, all of which are related to each other in some way. A particular verb sense—or multiple verb senses—may belong to many such groupings, resulting in overlapping collections of verb senses. Clustering has also been used to combine groupings with a high degree of overlap into more inclusive groupings.

4. Clustering Algorithm

Clustering is a technique that groups items together based on similarity between one or more of the features or attributes associated with the items. Here the items grouped together are verb senses, and lemmas used in their definitions are the attributes used for forming clusters. The assumption underlying that choice is that words that invoke a common semantic frame will tend to include the same concepts in their definitions, viz., those that refer to the corresponding situation and the participant structure associated with the frame. The restricted vocabulary used in the LDOCE definitions promotes the likelihood that when a given concept occurs in multiple definitions, the same lemma will be used to refer to the concept.

Clustering algorithms differ fundamentally in the similarity measures they use, the criteria they employ for using similarity values to form clusters, and the overall organization they impose on clusters. For this task, Voorhees' algorithm for the group average link method of hierarchical, agglomerative clustering has been used. Voorhees' (1986) algorithm uses the cosine coefficient as its similarity measure. In the context of the group average link method, which forms new clusters based on the average values of pairwise similarity values, the use of this similarity measure allows the algorithm to run in O(N) space, rather than the more typical (N^2) space, while continuing to run within (N^2) time.[2] Hierarchical, agglomerative clustering represents a choice among overall organizational structures. In a first pass every item (here, verb sense) is placed in its own cluster. In each subsequent pass the two most similar clusters are

[2]The space requirement drops to O(N) in the special case that between-item similarity (in our case, similarity between verb senses) is calculated as the inner product of appropriately weighted vectors. This works because cluster centroids, as the mean of all vectors in their spaces, can be used for computing between-cluster similarities (Voorhees 1986, p. 469).

```
/* Initialize */
        MaxSim  0;
        for (i  1 to CollectionSize) {
                create singleton cluster i for doc i; info[i].centroid  document i;
                ComputeSim (i, nn, sim);
                info[i].nn  nn; info[i].sim  sim; info [i].size  1;
                if (sim > MaxSim) {
                        id1  i; id2  nn; MaxSim  sim;
                }
        }

/* Merge clusters until only one cluster remains or remaining sims are 0 */
        while (MaxSim > 0.0 and NumActive > 1) {
                smaller  MIN (id1, id2); larger  MAX (id1, id2);
                info [smaller].centroid  MergeCentroids (smaller, larger);
                info [smaller].size  info [smaller].size + info [larger].size;
                a  index of larger in active;
                active [a]  active [NumActive]; NumActive  NumActive - 1;
                MergeClusters (smaller, larger, MaxSim);
                MaxSim  0;
                for (each cluster, a, in active) {
                        if (info [a].nn = larger or info [a].nn = smaller) {
                                FindMaxSim (a, nn, sim)  sim;
                                info [a].nn  nn; info [a].sim  sim;
                        }
                        if (info [a].sim > MaxSim) {
                                id1  a; id2  info [a].nn; MaxSim  info [a].sim;
                        }
                }
        }
```

Table 1. Group average link clustering method, Voorhees' algorithm

linked, forming a new cluster that replaces its input clusters. The clustering continues until either a single cluster is formed or a similarity threshold fails to be met. Voorhees' algorithm for constructing group average link hierarchies is outlined in table 1 (p. 471).

5. Clustering of Verb Senses Based on Words in Their Definitions

5.1 Pre-Processing

The input data underwent extensive pre-processing prior to the clustering process. Pre-processing included the following steps:

Phrasal verb senses (e.g., *strike out, take in*) were eliminated, leaving 12,663 LDOCE verb senses and 13,214 WordNet verb synsets (groups of synonymous word senses) for analysis.

After excluding usage notes, references to pictures, typographical codes, etc., from the LDOCE definitions, words from the definitions of LDOCE verb senses and WordNet verb synsets were extracted.

Words occurring on a stop word list were eliminated; the use of lists of 93 and 284 entries was explored.

Terms from definitions were stemmed, using the Porter (1980) stemmer. This stemmer is fairly conservative, only occasionally producing the same stem for word forms from different lemmas. However, it often fails to produce the same stem for two word forms from the same lemma. Given the relatively small size of the defining vocabulary used in LDOCE (some 2000+ words; see further discussion below), a list of equivalences between forms of words in the defining vocabulary, but not recognized by the Porter stemmer, has been developed by hand and applied to both the LDOCE and WordNet data sets.

Terms were assigned weights. A tf idf-like measure was used, with definitions taking the place of documents, except that if a term occurs in only one definition, its weight is set to an arbitrary low value, since it has no contribution to make to the clustering process. A further twist in the weighting scheme, affecting only a few terms in the processing of LDOCE, is explained below.

5.2 LDOCE

The *Longman Dictionary of Contemporary English* is intended primarily for non-native English speakers. To facilitate effective usage of the dictionary, a design decision was made to restrict the vocabulary of definitions and example sentences, insofar as it was possible, to a restricted set of basic English terms. This restricted vocabulary promotes the likelihood that when a given concept occurs in multiple definitions or example sentences, the same lemma will be used to refer to the concept. This, in turn, increases the probability that verb senses that invoke the same semantic frame will include the same terms in their definitions and thus may be identified through clustering.

Investigation of the restricted defining vocabulary has observed, however, that it does not altogether control for the standard lexical relationships to be found within a shared semantic frame. For example, all members of the following pairs (and triple) occur within the defining vocabulary: *difficult, hard; correct, right; evil, wicked; bind, fasten; cheat, deceive; look, see; say, speak, talk; different, other; nice, pleasant;* and *complete, all.* It is obvious that some senses of the members of these sets of words are quasi-synonymous; for clustering purposes, it would be best to treat them equivalently if synonymous senses are being used.

Establishing whether synonymous senses of these terms are in use is somewhat problematic. As basic English terms, the restricted defining vocabulary consists of high-frequency (Vossen et al., 1989, p. 174) and highly polysemous terms; Wilks et al. (1993, p. 349) show that the terms in the LDOCE restricted vocabulary have an average of twelve senses, which contrasts with an average of two senses for terms not in the defining vocabulary. However, LDOCE lexicographers were to use only 'central' senses of these terms. As a result, approximately one half were used in only one sense, another one quarter in two senses, and only one quarter in three or more senses (Wilks et al., 1993, p. 376).

Given this set of circumstances, I have developed the notion of a *strong sense* of a word, based on frequency data in SEMCOR, a WordNet-sense-tagged extract from the Brown Corpus. Where the use of a family of WordNet senses (those within the same part of speech and the same domain/lexicographic file) predominate over all other uses in SEMCOR, that family of senses is taken to be the strong sense of the word. Where such a strong sense occurs, it is assumed to be the sense used in LDOCE definitions.

Strong senses of words are then used to discover *strong relationships,* which occur where strong senses of nouns and verbs are linked in WordNet through hyponymy, hyperonymy, antonymy, entailment, and cause relationships. Strong relationships occur between the following sets of LDOCE defining terms:

be, exist, live
bind, fasten
bless, curse
borrow, lend
buy, give, pay, sell
cease, continue

fail, succeed, try
float, fly, sink, swim
forget, remember, think
hear, listen
know, recognize
laugh, weep
let, prevent
like, please
pronounce, say, speak, talk
ride, walk
shout, whisper
sleep, wake

The strong senses of the words in each of these sets are assumed to relate to a shared semantic frame. Therefore, each of the words in a strong relationship set is replaced in the verb sense data sets with a token common to the strong relationship (normally a string of all words in the strong relationship set separated by hyphens, e.g., *laugh-weep*). Weights for terms in strong relation sets were set to twice the average weight of the individual terms on the grounds that these terms had already undergone a rigorous grouping process and their grouping was deemed trustworthy.

The clustering algorithm has been executed numerous times, varying the stop word list used, the threshold for establishing strong relations, and the threshold used for establishing clusters. The output resulting from the combination of the smaller stop word list, the lesser strong relations threshold (predominance holds if half of all uses are attributable to a single family of senses), and a not-so-very-high/not-so-very-low cluster threshold criterion appears to yield the best output for the succeeding task, where reasonably high precision of the result is of greater import than high recall.

By way of example: On the one hand, on the basis of the occurrence of words such as 'actor', 'actress', 'part', 'play', and 'film' in their definitions, the clustering instantiation that generates the best output establishes relationships between the appropriate senses of *double, understudy, star,* and *premiere.* On the other hand, this clustering instantiation also establishes a relationship between senses of *latinize* and *pulsate* on the basis of the use of the word 'Latin' in both their definitions; however, the 'Latin' of 'Latin American dance music' in the definition of *pulsate* has nothing to do with the 'Latin language' of *latinize.* And because 'Latin' occurs only infrequently in the overall verb sense data set, the weighting it receives is fairly high. The incidence of such egregiously incorrect groupings appears to be felicitously low.

5.3 WordNet

WordNet is an online lexical database, widely used in the computational linguistics community for its broad coverage of general English and its relational organization. Synonymy and hyperonymy/hyponymy contribute significantly to the structure of the database. In particular, the basic structure in WordNet is the synset, a set of word senses that are (more or less) interchangeable in some circumstances. In the place of full definitions for its synsets, WordNet gives glosses designed to permit the native language speaker to differentiate between word senses. No particular attempt has been made to restrict the vocabulary used in the glosses.

WordNet data were pre-processed in the same general way as the LDOCE data, including the elimination of stop words from synset glosses, the stemming of remaining words, and the weighting of those stems. However, phrasal verbs were not eliminated, since synsets may include both phrasal and non-phrasal verbs. Nor was any attempt made to normalize the defining vocabulary of the glosses, as was done by applying equivalence sets based on strong relations to the LDOCE data. The clustering process used for the two data sets was the same, including the same threshold.

The results of clustering WordNet verb synsets can be surveyed through a group of computer-related clusters. The first of these clusters is comprised of the appropriate senses of the following verbs (while the clustering of WordNet data groups synsets together, these particular synsets mostly have but a single member): *initialize, hack, debug, download, upload, program, compile,* and *run*, based on occurrences of 'computer' and 'program' in their glosses. The second cluster includes *input, thrash, port, ftp, offload,* and *spool*, based on 'computer', 'transfer', and 'data' in their glosses. The occurrence of 'computer', 'disk', 'sector', 'record' and 'store' in their glosses bring together *write, format,* and *interleave*, while *cascade* and *close* are clustered together on the basis of 'computer', 'window', and 'desktop' in their glosses. A fifth, catch-all cluster associates together the appropriate senses of *emulate, swap, concatenate, scroll,* and *format*; here recurrence of the phrase 'computer science' accounts for the grouping. In considering these results, it should be recalled that it is not necessary that all of the "cue" words be found in every gloss. Also, because the clustering algorithm used produces a hierarchical output, a given synset will appear in at most one cluster. The hierarchical structure could be used to group more specific clusters into more general clusters, but the data have not yet been analyzed in that way.

6. Merging Verb Sense Clusters

Altogether, ten sources of data were manipulated to generate groups of LDOCE verb senses or WordNet verb synsets hypothesized to invoke the same semantic frame; terms in LDOCE verb sense definitions constitute one of those groups, while terms in the glosses of WordNet verb synsets constitute another. Each of the ten resulting groupings was processed to produce a list of pairs of semantically related LDOCE verb senses or WordNet verb synsets; WordNet verb synsets were then mapped to LDOCE verb senses. The combined data set consists of 25,063 distinctive pairs of semantically related LDOCE verb senses, accompanied by a count of how many data sources supported each pairing; all of these pairs meet either a weight or a count threshold. Of these pairs, 11,654 are supported by LDOCE data alone and 10,904 by WordNet data alone; only 2504 pairs are supported by both LDOCE and WordNet. Just over half the pairs are supported by a single data source; just under 50 are supported by 5 or more data sources.

The next step was to generate fully connected (i.e., complete link) verb sense groups. These groups are comprised of sets of verb senses, each of which was paired with every other verb sense in the group in the combined data set from the previous stage. There are 8,092 fully-connected verb sense groups, averaging 3.0 verb senses per group and including 7,098 distinct verb senses. That the number of groups and the number of verb senses in the groups are close in number reflects both the small size of the groups and especially the degree of overlap between groups. These sets of verb senses were then supplemented by the pairings from the LDOCE subject codes data source, a small (594) and almost perfect set of pairings. In this process, if one

To bogey1.1 is (in GOLF) to hit the ball into (a hole), taking one stroke more than is average.
To golf1.1 is to play golf. to go golfing.
To hit1.1 is to give a blow to; strike. He hit the other man. He hit the ball (with the BAT).
To hook1.4 is (of a ball) to travel in a HOOK. To hook1.4 is to hit (a ball) in a HOOK.
To loft1.1 is (in cricket and golf) to hit (a ball) high.
To mishit1.1 is (in cricket, GOLF, etc.) to hit (the ball) in a faulty way.
To par1.1 is (in the game of GOLF) to play the number of strokes for (a hole or all the holes) which is equal to PAR.
To pitch1.3 is (of a ball in cricket or GOLF) to hit the ground.
To pull1.11 is (in GOLF) to strike (the ball) to the left of the intended direction (or to the right if one is left-handed).
To putt1.1 is to strike (the ball) gently along the ground towards or into the hole.
To putt1.2 is to take a stated number of PUTTs to hit the ball into (the hole). Jacklin3-putted the 17th hole.
To shoot1.12 is (in GOLF) to make (the stated number of strokes) in playing a complete game. Miller shot a 69 today.
To sink1.10 is (in games like GOLF and BILLIARDS) to cause (a ball) to go into a hole.
To slice1.4 is to hit (a ball) in a SLICE.
To stymie1.1 is (in GOLF) cause (someone or oneself) to be stopped by positioning the balls in a STYMIE.
To top1.5 is (in GOLF) to hit (a ball) above the centre. He topped the ball and it went all along the ground.

Table 2. Glosses and example sentences in LDOCE for verb senses related to golf

verb sense of an LDOCE subject code pair appears in a verb sense group but the other does not, the second verb sense is added to the group.

The verb groups were then subjected to the same basic clustering process used for generating LDOCE verb sense clusters and WordNet verb synset clusters (no pre-processing was done, as it did not apply in these circumstances). In this context, verb groups were clustered together on the basis of verb senses occurring within them. Verb groups that clustered together in this process were then merged.

To illustrate this procedure, consider senses of verbs having to do with golf. At least 16 of these exist in LDOCE, as can be seen in table 2. One or more of these verb senses appears in 56 different fully connected verb sense groups (after supplementation by pairings from the LDOCE subject codes data source). As many as 8 of the verb senses belong to the same fully connected verb sense group, but many of the groups have only 2 members; the average membership size is 2.7. Clustering and merging reduces the number of golf-related clusters from 56 to 7. One of these clusters contains 14 of the 16 relevant verb senses.

7. Discussion

Of particular interest here is the effect of the restricted defining vocabulary in LDOCE on the success of the clustering venture. While the difference between the use of a semi-controlled vocabulary in the LDOCE definitions and the use of an open-ended vocabulary in the WordNet glosses is not the only difference encountered—as noted above, there are differences between intended users of the two resources, differences between definitions and glosses, and differences between using verb senses and using verb synsets as the basic unit of comparison—it is arguably the most important of those differences. First, it is the terms in the definitions/glosses on which the clustering is based. Second, differences in the defining vocabularies are simply a reflection of the difference in intended user (and to a lesser extent, of the difference between definitions and glosses). Third, the glosses given in WordNet could presumably be given separately and identically for each verb sense in a synset; in other words, a synset can be taken as a pre-clustering based on hypothetically identical glosses.

The table below summarizes the relevant statistical properties of the clustering process as applied to the word stems extracted from LDOCE definitions and from WordNet glosses.

Property	LDOCE	WordNet
Number of input units	12,663 (verb senses)	13,214 (synsets)
Number of distinct stems	3,292	6,344
Number of clusters generated	2,778	2,267
Number of input units represented in clustering output	8,084	5,092
Average number of members per cluster	2.9	2.25
Number of (LDOCE verb sense) pairs generated from clustering output	13,694	2,772

The number of verb *senses* in LDOCE and the number of verb *synsets* in WordNet are of equal magnitude. The number of verb *senses* in WordNet (24,169), however, is roughly double the number of its synsets, demonstrating that WordNet's coverage is more comprehensive and specific. The difference between LDOCE's restricted defining vocabulary and WordNet's open-ended defining vocabulary shows up clearly in the number of stems used in their vocabularies, with the WordNet vocabulary approximately twice the size of the LDOCE vocabulary. As a result of the less disparate input to the clustering process, LDOCE clustering produces a larger number of clusters than WordNet clustering does; the LDOCE clusters are on average larger as well. The proportion of input units retained within the output of the clustering process—64% for LDOCE, 39% for WordNet—also reflects the added power derived from the restricted defining vocabulary used in LDOCE. The huge difference in pairs generated by the two clusterings is only partially attributable to the larger number of clusters generated for LDOCE, the larger size of those clusters, and the larger number of distinct units (in this case, verb senses) retained in the clustering output. A significant part of this difference stems from repeated failures in the process of mapping WordNet verb senses to LDOCE verb senses to find equivalent senses. Some of these recall failures reflect differences in the defining vocabularies used, but many of them reflect gaps in LDOCE's coverage. There simply are many cases where WordNet verb senses have no counterpart in LDOCE.

8. Conclusion

Clustering has been used successfully within a research effort to identify semantic frames by grouping together LDOCE verb senses and WordNet verb synsets whose definitions/glosses use similar sets of stems. Clustering has also been used to merge similar sets of fully connected verb senses, thus significantly reducing redundancy in the data set. The use of a semi-controlled vocabulary for LDOCE definitions has provided a significant boost in the effort to generate extensional definitions of semantic frames.

Acknowledgements

Part of this work was funded by a National Science Foundation Graduate Research Fellowship.

References

Fillmore, Charles J. 1982. Frame semantics. In *Linguistics in the Morning Calm*, 111-137. Seoul: Hanshin.

Fillmore, Charles J. & Atkins, B. T. S. 1992. Towards a frame-based lexicon: The semantics of RISK and its neighbors. In A. Lehrer & E. F. Kittay (Eds.), *Frames, Fields, and Contrasts*, 75-102. Hillsdale, NJ: Erlbaum.

Porter, M.F. 1980. "An Algorithm for Suffix Stripping." *Program* 14/3: 130-137.

Voorhees, Ellen. 1986. Implementing agglomerative hierarchic clustering algorithms for use in document retrieval. *Information Processing & Management* 22/6: 465-476.

Vossen, Piek; Meijs, Willem; & den Broeder, Marianne. 1988. Meaning and structure in dictionary definitions. In B. Boguraev & T. Briscoe (Eds.), *Computational Lexicography for Natural Language Processing*, 171-192. London: Longman.

Wilks, Yorick; Fass, Dan; Guo, Cheng-Ming; McDonald, James; Plate, Tony; & Slator, Brian. 1993. Providing machine tractable dictionary tools. In J. Pustejovsky (Ed.), *Semantics and the Lexicon*, 341-401. Dordrecht, Boston, and London: Kluwer.

Discussion Points

11. What does the author mean by classification?

12. Does the research involve creation or implementation of a classification scheme?

13. How does the researcher use classification to improve the automated approach?

14. How do these methods compare to current human-generated approaches to classification?

15. How does the reported research expand our understanding of classification?

16. Does the research suggest an improvement over human-generated classification?

17. What do you think are the most important lessons learned in this research?

18. What do you think are the best practices reported in this research?

19. What would you recommend to the researcher as the next step in this approach?

20. Is there other related research that you would recommend the researcher become acquainted with?