Annotating WordNet

Helen Langone, Benjamin R. Haskell, George A. Miller
Cognitive Science Laboratory
Princeton University
{helen,ben,geo}@clarity.princeton.edu

Abstract

High-quality lexical resources are needed to both train and evaluate Word Sense Disambiguation (WSD) systems. The problem of ambiguity persists even in limited domains, thus the necessity for wide-coverage inventories of senses (dictionaries) and corpora sense-tagged to them. WordNet has been used extensively for WSD, for both its broad coverage and its large network of semantic relations. In this paper, we present a report on the state of our current endeavor to increase the connectivity of WordNet through sense-tagging the glosses, the result of which will be to create a more integrated lexical resource.

1 Introduction

High-quality lexical resources are needed to both train and evaluate Word Sense Disambiguation (WSD) systems. The problem of ambiguity persists even in limited domains, thus the necessity for wide-coverage inventories of senses (dictionaries) and corpora sense-tagged to them. WordNet (Miller et al., 1990; Fellbaum, ed., 1998) has been used extensively for WSD, both for its broad coverage and its large network of semantic relations. Entries in WordNet have, until now, been organized primarily around the semantic relations of synonymy, antonymy, hyponymy/troponymy, meronymy, and a few others which hold mainly among lexicalized concepts and word forms of the same grammatical class.[1] The noun and verb networks have been predominantly hierarchi-

cal, and the definitional glosses and illustrative sentences have not participated in the network of relations at all.

This paper reports on a project currently underway to sense-tag the glosses. Sense-tagging is the process of linking an instance of a word to the WordNet synset representing its context-appropriate meaning. Monosemous words[2] in the glosses can be tagged automatically, but in order to be truly reliable, the sense-tagging of polysemous words[3] must be done manually. This approach is in significant contrast with the work done at The University of Texas at Dallas on Extended WordNet[4], in which polysemous words in the WordNet glosses were sense-tagged primarily by automatic means. The result of the project described here will be to increase connectivity and make possible the association of words with related concepts that cut across grammatical class and hierarchy, providing a more integrated lexical resource.

2 Prior work

Previous efforts to sense-tag corpora manually have demonstrated that the task is not trivial. To begin with, the possibility of distinguishing word senses definitively in general is recognized as being problematic (Atkins and Levin, 1988; Kilgarriff, 1997; Hanks, 2000); indeed the notion of a "sense" has itself been the subject of long debate (see the Hanks and Kilgarriff papers for two recent contributions). These are topics in need of serious consideration, but outside of the scope of this paper. Certain issues emerge that are particularly relevant to designing

---

[1]The exceptions have been the links existing between adjectives derivationally related to nouns or verbs (*conceptual* is related to *concept* and *conceptuality*, *irritated* to *irritate*), and links between adverbs and the adjectives from which they derive (*absolutely* is related to WordNet sense 1 of *absolute*).

[2]Monosemous relative to WordNet (see footnote below).

[3]A word is polysemous if it has more than one related sense. A word with more than one *un*related sense is called a homonym. "Bank" is the classic example, its unrelated senses being "river bank" and "financial institution." WordNet does not make a distinction between homonymy and polysemy, therefore a monosemous word in WordNet is one which has neither related nor unrelated senses.

[4]http://xwn.hlt.utdallas.edu/

a manual sense-tagging system, and it is these we will concentrate on here.

The difficulties inherent in the sense-tagging task include the order in which words are presented to tag, a word's degree of polysemy and part of speech, vagueness of the context, the order in which senses are presented, granularity of the senses, and level of expertise of the person doing the tagging. Each will be addressed briefly in the following sections.

2.1 Targeted vs. sequential or 'lexical' vs. 'textual'

There are two approaches one can take to the order in which words are tagged. In the sequential approach, termed 'textual' by Kilgarriff (1998), the tagger proceeds through the text one word at a time, assigning the context-appropriate sense to each open class word as it is encountered. The targeted approach ('lexical' in Kilgarriff's terms) involves tagging all corpus instances of a pre-selected word, jumping through the text to each occurrence. The corpora produced by the SEMCOR and READER projects were tagged sequentially; the Kilo and HECTOR projects used the targeted approach (Kilo and READER are described in Miller et al. (1998), HECTOR in Atkins (1993), and SEMCOR in Fellbaum, ed. (1998)). In sequential tagging, the tagger is following the narrative, and so has the meaning of the text in mind when selecting senses for each new word encountered (context is foremost). In targeted tagging, the tagger has the various meanings of the word in mind when presented with each new context (sense distinctions are foremost). In their comparisons of the two approaches, Fellbaum et al. (2003) and Kilgarriff (1998) both conclude that sequential tagging increases the difficulty of the task by requiring the tagger to acquaint (and then reacquaint) themselves with the senses of a word each time they are confronted with it in the text. The targeted approach, on the other hand, enables the tagger to gain mastery of the sense-distinctions of a single word at a time, reducing the amount of effort required to tag each new instance. Miller et al. (1998) present a contrasting view. In evaluating the Kilo and READER tagging tasks, they find targeted tagging to be more tedious for the taggers than sequential tagging, and no faster, as time is needed to assimilate new contexts for each word occurrence.

2.2 Polysemy, POS, and sense order

In their 1997 paper, Fellbaum et al. analyzed the results of the SEMCOR project, in which part of the Brown Corpus (Kučera and Francis, 1967) was tagged to WordNet 1.4 senses. Their analysis identified three factors that influenced the difficulty level, and thus the accuracy, of the tagging task: degree of polysemy of the word being tagged, the word's part of speech, and the order in which the WordNet sense choices are displayed to the person

doing the tagging.

The effect of a high degree of polysemy is to present more choices to the tagger, usually with finer distinctions among the senses, increasing the difficulty of selecting one out of several closely-related senses.

The correspondence of a word's part of speech with accuracy of tagging stems from the nature of the objects that words of a certain class denote. Words that refer to concepts that are concrete tend to have relatively fixed, easily distinguishable senses. Words with more abstract or generic referents tend to have a more flexible semantics, with meanings being assigned in context, and hence more difficult to pin down. Nouns tend to be in the former category, verbs in the latter. More abstract classes also tend to have a higher degree of polysemy, adding to the effect.

Finally, the presentation of the sense choices in WordNet order, with the most frequent sense first, creates a bias towards selecting the first sense. Their study shows that randomly ordering the senses removes this effect.

2.3 Granularity of senses

Palmer et al. (to appear, 2004) examines the relationship between manual and automatic tagging accuracy and the granularity of the sense inventory. Granularity has to do with fineness of distinctions made from a lexicographer's point of view, and not the number of senses that a word form exhibits in context. It is related to polysemy, in that the greater a word's degree of polysemy, the finer the distinctions that can be made in defining senses for that word. In their experiment, Palmer et al. have lexicographers group WordNet 1.7[5] senses according to syntactic and semantic criteria, which are used by taggers to tag corpus instances. An automatic WSD system trained on the data tagged using grouped senses shows a 10% overall improvement in performance against running it on data tagged without using the groupings. Their study shows that improvement came *not* from the fewer number of senses resulting from the groupings, but from the groupings themselves, which increased the manual tags' accuracy (defined as agreement between taggers), thereby increasing the accuracy of the systems that learned from them. This effect arises from the slippery nature of word senses and the impossibility of capturing them in neatly delimited, universally agreed-upon sense-boxes. New usages of words extending old meanings, vague contexts that select for multiple senses, and the limits of the tagger's own knowledge of a specialized domain, all defy the assignment of a single, unequivocal sense to a word's instance across annotators. Palmer et al. propose sense groupings as a practical solution in these situations.

---

[5]The words used for this experiment were the polysemous verbs from the lexical sample task for SENSEVAL-2 (Edmonds and Cotton, 2001).

2.4 Tagger expertise

Finally, there is the question of whether novice taggers with adequate training can attain the level of accuracy of experienced lexicographers and linguists. Fellbaum et al. (1997) answer this in the negative. Their findings show novice tagger accuracy decreasing as the number of senses, or fineness of distinctions among the senses, increases. Level of expertise likely influenced the slow pace of tagging reported for the Kilo and READER projects, which employed novice taggers. During the tagging of the evaluation dataset for SENSEVAL-1, the highly experienced lexicographers who did the tagging reported the time spent absorbing new contexts dropped off rapidly after a slow start-up period (Krishnamurthy and Nicholls, 2000).

2.5 The present approach

To the extent that these difficulties can be addressed, we have attempted to do so. We feel the most accurate results can be obtained from the targeted approach using linguistically-trained taggers. The nature of the glosses (relatively short, completely self-contained) means that a fairly restricted context will need to be assimilated for each instance of a token, eliminating one factor of difficulty associated with the targeted approach. Since a definition is, by definition, unambiguous, the context provided by a gloss should, in theory, never be insufficient to disambiguate the words used within it. In this respect, the glosses differ from KWIC (Key Word In Context) lines in a concordance, with which they can be compared. KWIC concordances, so named because they display corpus instances of a (key) word along with surrounding text, are used by lexicographers in a manner very similar to the targeted approach to sense-tagging. There the task is to define the word, to determine and delineate its senses given its contexts of use. One further difference to be exploited is the fact that, unlike a sentence in a typical corpus, a gloss is embedded within a network of WordNet relations. This means that immediate hypernym, domain category, and other relations can be made available to the user as additional disambiguation aids.

The order of the senses will be scrambled in the manual tagging interface so as to prevent a bias towards the first sense listed. To avoid putting any additional burden on the tagger, the order of senses will be fixed at the beginning of the session, and kept constant until the tagger exits the program or selects another word to tag.

Underspecified word senses are expressed in WordNet in the form of verb groups[6], which will be presented to the tagger in the sense display with the option to select either the entire group, or individual senses within the

group. Where no appropriate grouping exists, and context and domain category are not enough to fully disambiguate, multiple tags can be assigned. Precise guidelines for when multiple senses can be assigned, and under what criteria, will need to be developed, and taggers will need to be extensively trained on them.

3 Annotating the glosses

There are six major components to the present sense-tagging system. They function in pipeline fashion, with the output of one being fed as input to the next. Each pass through the data produces output in valid XML, the structure of which is covered in Subsection 3.1. The six components are: (1) a gloss parser, (2) a tokenizer, (3) an ignorable text chunker and classifier, (4) a WordNet collocation[7] recognizer, (5) an automatic sense tagger, and (6) a manual tagging tool.

Prior to and in conjunction with building the preprocessor (the first four components), analysis of the WordNet glosses was undertaken to determine what should be presented for tagging, and what was not to be tagged. Ignorable classes of word forms and multi-word forms were determined during this phase. These were used as a basis for the development of a stop list of words and phrases to ignore completely, and a second, semi-stop list that we have dubbed the "wait list". The stop list is reserved for unequivocally closed-class words and phrases including prepositions, conjunctions, determiners, pronouns and modals, plus multi-word forms that function as prepositions (e.g., "by means of"). Words on the wait list will be held out from the automatic tagging stage for manual review and tagging later. Since WordNet covers only open-class parts of speech, word forms that have homographs in both open and closed-class parts of speech are on this list. During the manual tagging stage, the open-class senses will be tagged. Highly polysemous words such as "be" and "have" are also waitlisted.

Many glosses also contain example sentences. While not an essential part of the semantic makeup of the synset, they do give some information about the illustrated word's sense-specific context of use, contributing to meaning in a different way.[8] For this reason, we will be tagging the synset word (and only that word) of which the sentence is an exemplar. By virtue of being located *within* the synset, the exemplified form is in effect automatically disambiguated—it's just a matter of assigning the tag.

---

[6]Groupings exist for many, though not all, polysemous verbs in WordNet.

[7]In WordNet terminology, a collocation is a superordinate term for a variety of multi-word forms, including, but not restricted to, names, compounds, phrasal verbs, and idiomatic phrases.

[8]Insofar as meaning is defined in part by use.

3.1 The glosstag DTD

Development of the formal model for the sense-tagged glosses took the DTD from the SEMCOR project as a starting point (Landes et al., 1998). It went through several iterations of modification, first to accommodate the specifics of the dataset being tagged (WordNet glosses as opposed to open text), and then to refine the handling of WordNet collocations. Prior tagging efforts had employed the WordNet method of representing collocations as single word forms, with underscores replacing the spaces between words. While it is a practical solution that gives the collocation the same status and representational form that it has as an entry in WordNet, by treating a collocation as a "word", we lose the fact that it is decomposable into smaller units. This renders difficult the coding of discontinuous collocations (that is, collocations interrupted by one or more intervening words, for example 'His performance blew the competition out of the water', where "blow out of the water" is a WordNet collocation). A scheme that enables collocations to be treated both as individual words and as multi-word units is therefore desirable, particularly if future parsing passes need to identify the internal structure of a collocation, as for distinguishing phrase heads from non-heads.

The smallest structural unit, then, is a word (or piece of punctuation), marked as <cf> if it is part of a WordNet collocation, and as <wf> otherwise. Attributes on the <wf> and <cf> elements identify each form uniquely in the gloss, and link together the constituent <cf>'s of a collocation.

The major structural units of a gloss are <def>, <ex>, and <aux>. <def> contains the definitional portion of the gloss, the main interest of the tagging task. A <def> may be followed by one or more <ex>'s, each containing an example sentence. Auxiliary information,[9] coded as <aux>, may precede or follow the <def>, or occur within it. Figure 1 shows the marked up gloss for sense 11 of **life** (the gloss text is "living things collectively;"), as it looks after preprocessing.

Prior to sense-tagging, the lemma attribute on the <wf> or head[10] <cf> of a collocation is set to all possible lemma forms, as determined during lemmatization (explicated more fully below). After sense-tagging, the lemma attribute is set to only the lemma of the word/collocation that it is tagged to, all other options are deleted. An <id> element representing the sense tag is inserted as a child of the <wf> (or <cf>), if multi-

ple sense-tags are assigned, then multiple <id>'s are assigned, one for each tag. Figure 2 shows the sense-tagged gloss for **life**.

3.2 Preprocessing and automatic tagging

The preprocessing stage segments the gloss into chunks and tokenizes the gloss contents into words and WordNet collocations. The tokenization pass isolates word forms and disambiguates lexical from non-lexical punctuation. Lexical punctuation is retained as part of the word form, non-lexical punctuation is encoded as ignorable <wf>'s. Abbreviations and acronyms are recognized, contractions split, and stop list and wait list forms are handled. All <wf>'s other than punctuation are lemmatized, that is, they are reduced to their WordNet entry form using an in-house tool, moan[11], that was developed for this purpose. Part of speech is not disambiguated during preprocessing, therefore lemmatizing assigns all potential lemma forms for all part of speech classes that moan returns for the token. Part of speech disambiguation will occur as a side-effect of sense-tagging, avoiding the introduction of errors related to POS-tagging. Lemmatizing serves two functions, first when searching the database of glosses for the term being tagged, and then when displaying the sense choices for a particular instance.

The targeted tagging approach introduces the problem of locating all inflected forms of the word/collocation to be tagged. Rather than build a tool to generate inflected forms, our solution was to pre-lemmatize the corpus and search on the lemma forms, on the assumption that while the search will overgenerate matches, it will not miss any. Locating alternations in hyphenation will be handled in a similar way, via the pre-indexing of alternate forms of hyphenated words/collocations in WordNet.

The ignorable text classifier recognizes ignorable text as described earlier, chunking multi-word terms and assigning attributes indicating semantic class. The markup will enable them to be treated as individual words or, alternatively, as a single form indicating the class, which will be of use should further parsing or semantic processing of the glosses be called for.

The WordNet collocation recognizer, or globber, uses a bag-of-words approach to locate multi-word forms in the glosses. First, all possible collocations are pulled from the WordNet database. This list is then filtered by several criteria designed to exclude candidates that cannot be accurately identified automatically. The largest class of excluded words is that of phrasal verbs, which cannot

---

[9]Auxiliary text is a cover term for a range of numeric and symbolic classes (dates, times, numbers, numeric ranges, measurements, formulas, and the like), and parenthesized and other secondary text that are inessential to the meaning of a synset.

[10]"Head" here refers simply to the first word in the collocation, and not the syntactic head. The head <cf> bears the lemma and sense-tag(s) for the entire collocation.

[11]Moan falls somewhere between a stemmer and full morphological analyzer—it recognizes inflectional endings and restores a corpus instance to its possible lemma form(s) classified by part of speech and grammatical category of the inflectional suffix. Lemma form is the WordNet entry spelling, if the word is in WordNet.

```
<synset pos="n" ofs="00005905">
  <gloss desc="wsd">
    <def>
      <wf wf-num="1" tag="un" lemma="living%1|live%2|living%3">living</wf>
      <wf wf-num="2" tag="un" lemma="thing%1|things%1">things</wf>
      <wf wf-num="3" tag="un" lemma="collectively%4" sep="">collectively</wf>
      <wf wf-num="4" type="punc" tag="ignore">;</wf>
    </def>
  </gloss>
</synset>
```

Figure 1: Preprocessed gloss for **life**, prior to semantic annotation

```
<synset pos="n" ofs="00005905">
  <gloss desc="wsd">
    <def>
      <cf coll="a" tag="auto">
        <glob coll="a" tag="auto">
          <id coll="a" lemma="living_thing" pos="n" ofs="00004323"/>
        </glob>
        living
      </cf>
      <cf coll="a" tag="cf">things</cf>
      <wf tag="auto" sep="">
        <id lemma="collectively" pos="r" ofs="00119700"/>
        collectively
      </wf>
      <wf tag="ignore" type="punc">;</wf>
    </def>
  </gloss>
</synset>
```

Figure 2: Semantically-annotated gloss for **life**

easily be distinguished from verbs followed by prepositions heading prepositional phrases.[12] Many of these will be globbed by hand in the early stages of manual tagging.

From this list of excluded words, we also generate a list of collocations that contain monosemous words. This list will later be used to prevent those words from being erroneously tagged in the automatic sense-tagging stage. The final list of words to be automatically globbed also takes into account variations in hyphenation and capitalization.

Once the list is completed, the next step is to create an index of the glosses referenced by the lemmatized forms they contain. For each collocation, the globber calculates the intersection of the lists of glosses containing its constituent word forms. This list of possible collocations is then ordered by gloss.

The final step of the globber iterates through each of the glosses, three passes per gloss. The first pass marks the monosemous words found in excluded collocations, without globbing the collocation. Pass two identifies multi-word forms that appear as consecutive <wf>'s in the text. The final pass attempts to locate disjoint collocations that follow certain set patterns of usage, such as "ribbon, bull, and garter snakes", where "ribbon snake", "bull snake", and "garter snake" are all in WordNet. "Garter snake" is globbed in pass two, and parallel structure helps identify "ribbon snake" and "bull snake" in the third pass.

After preprocessing is complete, the automatic sense tagger tags monosemous <wf>'s and <cf>'s to their WordNet senses. Words and collocations tagged by the automatic tagger are distinguished from manually tagged terms by an attribute in the markup.

Sense-tagging the glosses to WordNet senses presupposes that all words used in the glosses (and all senses used of those words) exist as entries. The preprocessing and auto-tagging phase will therefore include a few dry runs to identify any typographical errors and words not covered, errors will be fixed and open class words or word senses will be added to WordNet as necessary.

3.3 Manual tagger interface

The single most important design consideration for the manual tagger interface is the repetitiveness inherent to the task. With approximately 550,000 polysemous open-class words and collocations[13] in the glosses, each tagger will tag hundreds of words in a day of work. We have made every effort to minimize the amount of mouse movement and the number of button presses required to tag each word.

The layout of the program window is simple. The current search term is displayed in an entry box near the top of the screen. Below this box are two text boxes, for glosses and examples, respectively. Buttons used to alter the current tag or tags lie above the final text box, which is used to display and select the WordNet sense or senses for the current word.

The tag status of each word or collocation in the gloss and example boxes is indicated through the use of color, font, and highlighting. Orange text indicates a term that has been automatically tagged, red type denotes a manually tagged word, words marked as ignorable are shown in black, and the remainder of the taggable text is blue. Words that are part of a collocation are underlined, and forms that match the targeted search term are bolded. The current selection is highlighted in yellow.

There are several ways to navigate the glosses. For targeted tagging, the user chooses one or more senses, then clicks the 'Tag' button, assigning those senses and automatically jumping to the next untagged instance of the search word. Other buttons allow movement to the next or previous instance of the search term without altering the tag status of the current selection. The interface was designed with targeted tagging in mind, but the user can switch between targeted and sequential modes, to fix an improperly tagged word.

The interface allows a user to filter the displayed senses by part of speech, to concentrate on the relevant options. When context is insufficient to fully disambiguate, a word or collocation can be tagged to more than one sense or to a WordNet verb group. To prevent bias caused by the order of the displayed senses, each time a new targeted search term is entered, the senses shown in the sense box are shuffled after being grouped by part of speech.

During the targeted tagging process, the interface also enables a user to easily inspect and change the sense tags assigned to words other than the search term. The interface will display a box containing a tagged word's senses when the cursor is placed over it, providing useful information for disambiguating the search term. Additionally, if the user notices a tagging error, the mis-tagged word can be selected for editing. Errors and omissions of globbing can be corrected in a similar fashion. To "un-glob" a collocation, one need only click on the collocation and click on the "un-glob" button. To group separate <wf>s into a new collocation, a user can select each constituent form with the mouse and click the "glob" button. The interface will then provide a list of potential lemmas for the collocation, from which the user can select the appropriate choice.

4 Looking ahead

WordNet currently consists of 146,000 lexemes organized into more than 116,000 synsets. There are over 117,000 bidirectional links comprising the hyponymy,

---

[12]"Last year legal fees ate up our profits." versus "Last night we ate up the street."

[13]With an average polysemy of 2.59 senses per word form.

troponymy, and meronymy hierarchies, and 4,000 bidirectional antonym links. Once completed, the sense-tagged glosses will contribute an estimated 800,000 links to the network, increasing the internal connectivity and associating words with related concepts that cut across grammatical class and hierarchy. From the synset for a word, the synsets for all conceptually-related words used in its gloss can be accessed via their sense tags. From *those* synsets, hierarchical and other semantic links can be followed, as well the sense tags in those glosses, situating the word in an ever-expanding network of links. The sense-tagged glosses will provide an additional dimension of meaning not expressible by purely hierarchical relations. The hierarchical structure of WordNet represents sense distinctions that stem from a polysemous word's distinct superordinates (paradigmatic difference). Not represented are a word's syntagmatic properties—the ways in which meaning is constrained by the differing contexts in which a polysemous word appears. The sense-tagged glosses, taken as a corpus of disambiguated contexts for a word, provide just that.

NLP needs high-quality sense-tagged corpora and sense inventories. The project currently underway is a step towards providing both in one integrated lexical resource.

References

Sue Atkins. 1993. Tools for computer-aided lexicography: the Hector project. *Papers in Computational Lexicography: COMPLEX '93.* Budapest.

Beryl Atkins and Beth Levin. 1988. Admitting impediments. *Proceedings of the 4th Annual Conference of the UW Center for the New OED.* Oxford, UK.

Philip Edmonds and Scott Cotton. 2001. SENSEVAL-2: Overview. *Proceedings of SENSEVAL-2: Second International Workshop on Evaluating Word Sense Disambiguation Systems.* Toulouse, France.

Christiane Fellbaum, ed. 1998. *WordNet: An Electronic Lexical Database.* MIT Press, Cambridge, MA.

Christiane Fellbaum, Lauren Delfs, Susanne Wolff, and Martha Palmer. 2003. Word meaning in dictionaries, corpora, and the speaker's mind. In *Meaningful texts: The Extraction of Semantic Information from Monolingual and Multilingual Corpora*, eds. Geoff Barnbrook, Pernilla Danielsson and Michaela Mahlberg. Birmingham University Press, Birmingham, UK.

Christiane Fellbaum, Joachim Grabowski, and Shari Landes. 1997. Analysis of a hand-tagging task. *Proceedings of the ACL/Siglex workshop.* Somerset, NJ.

Patrick Hanks. 2000. Do word meanings exist? *Computers and the Humanities*, 34(1-2):205-215. Special Issue on SENSEVAL.

Adam Kilgarriff. 1997. I don't believe in word senses. *Computers and the Humanities*, 31(2):91-113.

Adam Kilgarriff. 1998. Gold standard datasets for evaluating Word Sense Disambiguation programs. *Computer Speech and Language*, 12(3):453-472.

Ramesh Krishnamurthy and Diane Nicholls. 2000. Peeling an onion: the lexicographer's experience of manual sense-tagging. *Computers and the Humanities*, 34(1-2):85-97. Special Issue on SENSEVAL.

Henry Kučera and W. Nelson Francis. 1967. *The standard corpus of present-day American English.* (Electronic database.) Brown University, Providence, RI.

Shari Landes, Claudia Leacock, and Randee I. Tengi. 1998. Building semantic concordances. In *WordNet: An Electronic Lexical Database*, ed. Christiane Fellbaum. MIT Press, Cambridge, MA.

George A. Miller, Richard Beckwith, Christiane Fellbaum, Derek Gross, and Katherine J. Miller. 1990. Introduction to WordNet: an on-line lexical database. *International Journal of Lexicography*, 3(4):235-244.

George A. Miller, Randee Tengi, and Shari Landes. 1998. Matching the tagging to the task. In *WordNet: An Electronic Lexical Database*, ed. Christiane Fellbaum. MIT Press, Cambridge, MA.

Martha Palmer, Hoa Trang Dang, and Christiane Fellbaum. To appear, 2004. Making fine-grained and coarse-grained sense distinctions, both manually and automatically. *Natural Language Engineering.*