Role of Verbs in Document Analysis

Judith Klavans* and Min-Yen Kan**

Center for Research on Information Access* and Department of Computer Science**
Columbia University
New York, NY 10027, USA

Abstract

We present results of two methods for assessing the event profile of news articles as a function of verb type. The unique contribution of this research is the focus on the role of verbs, rather than nouns. Two algorithms are presented and evaluated, one of which is shown to accurately discriminate documents by type and semantic properties, i.e. the event profile. The initial method, using WordNet (Miller et al. 1990), produced multiple cross-classification of articles, primarily due to the bushy nature of the verb tree coupled with the sense disambiguation problem. Our second approach using English Verb Classes and Alternations (EVCA) Levin (1993) showed that monosemous categorization of the frequent verbs in WSJ made it possible to usefully discriminate documents. For example, our results show that articles in which communication verbs predominate tend to be opinion pieces, whereas articles with a high percentage of agreement verbs tend to be about mergers or legal cases. An evaluation is performed on the results using Kendall's τ . We present convincing evidence for using verb semantic classes as a discriminant in document classification.¹

1 Motivation

We present techniques to characterize document type and event by using semantic classification of verbs. The intuition motivating our research is illustrated by an examination of the role of

nouns and verbs in documents. The listing below shows the ontological categories which express the fundamental conceptual components of propositions, using the framework of Jackendoff (1983). Each category permits the formation of a wh-question, e.g. for [THING] "what did you buy?" can be answered by the noun "a fish". The wh-questions for [ACTION] and [EVENT] can only be answered by verbal constructions, e.g. in the question "what did you do?", where the response must be a verb, e.g. *jog, write, fall*, etc.

[THING]	[DIRECTION]	[ACTION]
[PLACE]	[MANNER]	[EVENT]
[AMOUNT]		

The distinction in the ontological categories of nouns and verbs is reflected in information extraction systems. For example, given the noun phrases *fares* and *US Air* that occur within a particular article, the reader will know what the story is about, i.e. *fares* and *US Air*. However, the reader will not know the [EVENT], i.e. what happened to the *fares* or to *US Air*. Did airfare prices *rise, fall* or *stabilize*? These are the verbs most typically applicable to prices, and which embody the event.

1.1 Focus on the Noun

Many natural language analysis systems focus on nouns and noun phrases in order to identify information on who, what, and where. For example, in summarization, Barzilay and Elhadad (1997) and Lin and Hovy (1997) focus on multi-word noun phrases. For information extraction tasks, such as the DARPA-sponsored Message Understanding Conferences (1992), only a few projects use verb phrases (events), e.g. Appelt et al. (1993), Lin (1993). In contrast, the named entity task, which identifies nouns and noun phrases, has generated numerous projects

¹The authors acknowledge earlier implementations by James Shaw, and very valuable discussion from Vasileios Hatzivassiloglou, Kathleen McKeown and Nina Wacholder. Partial funding for this project was provided by NSF award #IRI-9618797 STIMULATE: Generating Coherent Summaries of On-Line Documents: Combining Statistical and Symbolic Techniques (co-PI's McKeown and Klavans), and by the Columbia University Center for Research on Information Access.

as evidenced by a host of papers in recent conferences, (e.g. Wacholder et al. 1997, Palmer and Day 1997, Neumann et al. 1997). Although rich information on nominal participants, actors, and other entities is provided, the named entity task provides no information on **what happened** in the document, i.e. the **event** or **action**. Less progress has been made on ways to utilize verbal information efficiently. In earlier systems with stemming, many of the verbal and nominal forms were conflated, sometimes erroneously. With the development of more sophisticated tools, such as part of speech taggers, more accurate verb phrase identification is possible. We present in this paper an effective way to utilize verbal information for document type discrimination.

1.2 Focus on the Verb

Our initial observations suggested that both occurrence and distribution of verbs in news articles provide meaningful insights into both article type and content. Exploratory analysis of parsed Wall Street Journal data² suggested that articles characterized by movement verbs such as *drop*, *plunge*, or *fall* have a different event profile from articles with a high percentage of communication verbs, such as *report*, *say*, *comment*, or *complain*. However, without associated nominal arguments, it is impossible to know whether the [THING] that *drops* refers to airfare prices or projected earnings.

In this paper, we assume that the set of verbs in a document, when considered as a whole, can be viewed as part of the conceptual map of the events and action in a document, in the same way that the set of nouns has been used as a concept map for entities. This paper reports on two methods using verbs to determine an event profile of the document, while also reliably categorizing documents by type. Intuitively, the event profile refers to the classification of an article by the kind of event. For example, the article could be a discussion event, a reporting event, or an argument event.

To illustrate, consider a sample article from WSJ of average length (12 sentences in length) with a high percentage of communication verbs. The profile of the article shows that there are 19 verbs: 11 (57%) are communication verbs, including *add*, *report*, *say*, and *tell*. Other

verbs include *be skeptical*, *carry*, *produce*, and *close*. Representative nouns include *Polaroid Corp.*, *Michael Ellmann*, *Wertheim Schroder & Co.*, *Prudential-Bache*, *savings*, *operating results*, *gain*, *revenue*, *cuts*, *profit*, *loss*, *sales*, *analyst*, and *spokesman*.

In this case, the verbs clearly contribute information that this article is a report with more opinions than new facts. The preponderance of communication verbs, coupled with proper noun subjects and human nouns (e.g. *spokesman*, *analyst*) suggest a discussion article. If verbs are ignored, this fact would be overlooked. Matches on frequent nouns like *gain* and *loss* do not discriminate this article from one which announces a *gain* or *loss* as breaking news; indeed, according to our results, a breaking news article would feature a higher percentage of motion verbs rather than verbs of communication.

1.3 On Genre Detection

Verbs are an important factor in providing an event profile, which in turn might be used in categorizing articles into different genres. Turning to the literature in genre classification, Biber (1989) outlines five dimensions which can be used to characterize genre. Properties for distinguishing dimensions include verbal features such as tense, agentless passives and infinitives. Biber also refers to three verb classes: private, public, and suasive verbs. Karlgren and Cutting (1994) take a computationally tractable set of these properties and use them to compute a score to recognize text genre using discriminant analysis. The only verbal feature used in their study is present-tense verb count. As Karlgren and Cutting show, their techniques are effective in genre categorization, but they do not claim to show how genres differ. Kessler et al. (1997) discuss some of the complexities in automatic detection of genre using a set of computationally efficient cues, such as punctuation, abbreviations, or presence of Latinate suffixes. The taxonomy of genres and facets developed in Kessler et al. is useful for a wide range of types, such as found in the Brown corpus. Although some of their discriminators could be useful for news articles (e.g. presence of second person pronoun tends to indicate a letter to the editor), the indicators do not appear to be directly applicable to a finer classification of news articles.

News articles can be divided into several stan-

²Penn TreeBank (Marcus et al. 1994) from the Linguistic Data Consortium.

dard categories typically addressed in journalism textbooks. We base our article category ontology, shown in lowercase, on Hill and Breen (1977), in uppercase:

1. FEATURE STORIES : feature;
2. INTERPRETIVE STORIES: editorial, opinion, report;
3. PROFILES;
4. PRESS RELEASES: announcements, mergers, legal cases;
5. OBITUARIES;
6. STATISTICAL INTERPRETATION: posted earnings;
7. ANECDOTES;
8. OTHER: poems.

The goal of our research is to identify the role of verbs, keeping in mind that event profile is but one of many factors in determining text type. In our study, we explored the contribution of verbs as one factor in document type discrimination; we show how article types can be successfully classified within the news domain using verb semantic classes.

2 Initial Observations

We initially considered two specific categories of verbs in the corpus: communication verbs and support verbs. In the WSJ corpus, the two most common main verbs are *say*, a communication verb, and *be*, a support verb. In addition to *say*, other high frequency communication verbs include *report*, *announce*, and *state*. In journalistic prose, as seen by the statistics in Table 1, at least 20% of the sentences contain communication verbs such as *say* and *announce*; these sentences report point of view or indicate an attributed comment. In these cases, the subordinated complement represents the main event, e.g. in “Advisors *announced* that IBM stock *rose* 36 points over a three year period,” there are two actions: *announce* and *rise*. In sentences with a communication verb as main verb we considered both the main and the subordinate verb; this decision augmented our verb count an additional 20% and, even more importantly, further captured information on the actual event in an article, not just the communication event. As shown in Table 1, support verbs, such as *go* (“go out of business”) or *get* (“get along”), constitute 30%, and other content verbs, such as *fall*, *adapt*, *recognize*, or *vow*, make up the remaining 50%. If we exclude all support type verbs, 70% of the verbs yield information in answering the question “what happened?” or “what did X do?”

3 Event Profile: WordNet and EVCA

Since our first intuition of the data suggested that articles with a preponderance of verbs of

Verb Type	Sample Verbs	%
communication	say, announce, ...	20%
support	have, get, go, ...	30%
remainder	abuse, claim, offer, ...	50%

Table 1: Approximate Frequency of verbs by type from the *Wall Street Journal* (main and selected subordinate verbs, $n = 10,295$).

a certain semantic type might reveal aspects of document type, we tested the hypothesis that verbs could be used as a predictor in providing an event profile. We developed two algorithms to: (1) explore WordNet (WN-Verber) to cluster related verbs and build a set of verb chains in a document, much as Morris and Hirst (1991) used Roget’s Thesaurus or like Hirst and St. Onge (1998) used WordNet to build noun chains; (2) classify verbs according to a semantic classification system, in this case, using Levin’s (1993) *English Verb Classes and Alternations* (EVCA-Verber) as a basis. For source material, we used the manually-parsed Linguistic Data Consortium’s *Wall Street Journal* (WSJ) corpus from which we extracted main and complement of communication verbs to test the algorithms on.

Using WordNet. Our first technique was to use WordNet to build links between verbs and to provide a semantic profile of the document. WordNet is a general lexical resource in which words are organized into synonym sets, each representing one underlying lexical concept (Miller et al. 1990). These synonym sets – or synsets – are connected by different semantic relationships such as hypernymy (i.e. *plunging* is a way of *descending*), synonymy, antonymy, and others (see Fellbaum 1990). The determination of relatedness via taxonomic relations has a rich history (see Resnik 1993 for a review). The premise is that words with similar meanings will be located relatively close to each other in the hierarchy. Figure 1 shows the verbs *cite* and *post*, which are related via a common ancestor *inform*, ..., *let know*.

The WN-Verber tool. We used the hypernym relationship in WordNet because of its high coverage. We counted the number of edges needed to find a common ancestor for a pair of verbs. Given the hierarchical structure of WordNet, the lower the edge count, in principle, the closer the verbs are semantically. Because WordNet

Figure 1: Taxonomic Relations for *cite* and *post* in WordNet.

allows individual words (via synsets) to be the descendent of possibly more than one ancestor, two words can often be related by more than one common ancestor via different paths, possibly with the same relationship (grandparent and grandparent, or with different relations (grandparent and uncle).

Results from WN-Verber. We ran all articles longer than 10 sentences in the WSJ corpus (1236 articles) through WN-Verber. Output showed that several verbs – e.g. *go*, *take*, and *say* – participate in a very large percentage of the high frequency synsets (approximate 30%). This is due to the width of the verb forest in WordNet (see Fellbaum 1990); top level verb synsets tend to have a large number of descendants which are arranged in fewer generations, resulting in a flat and bushy tree structure. For example, a top level verb synset, *inform, ..., give information, let know* has over 40 children, whereas a similar top level noun synset, *entity*, only has 15 children. As a result, using fewer than two levels resulted in groupings that were too limited to aggregate verbs effectively. Thus, for our system, we allowed up to two edges to intervene between a common ancestor synset and each of the verbs’ respective synsets, as in Figure 2.

Figure 2: Configurations for relating verbs in our system.

In addition to the problem of the flat nature of the verb hierarchy, our results from WN-Verber are degraded by ambiguity; similar effects have been reported for nouns. Verbs with differences in high versus low frequency senses caused certain verbs to be incorrectly related;

for example, *have* and *drop* are related by the synset meaning “to give birth” although this sense of *drop* is rare in WSJ.

The results of WN-Verber in Table 2 reflect the effects of bushiness and ambiguity. The five most frequent synsets are given in column 1; column 2 shows some typical verbs which participate in the clustering; column 3 shows the type of article which tends to contain these synsets. Most articles ($864/1236 = 70\%$) end up in the top five nodes. This illustrates the ineffectiveness of these most frequent WordNet synset to discriminate between article types.

Synset	Sample Verbs in Synset	Article types (listed in order)
Act (interact, act together, ...)	have, relate, give, tell	announcements, editorials, features
Communicate (communicate, intercommunicate, ...)	give, get, inform, tell	announcements, editorials, features, poems
Change (change)	have, modify, take	poems, editorials, announcements, features
Alter (alter, change)	convert, make, get	announcements, poems, editorials
Inform (inform, round on, ...)	inform, explain, describe	announcements, poems, features

Table 2: Frequent synsets and article types.

Evaluation using Kendall’s Tau. We sought independent confirmation to assess the correlation between two variables’ rank for WN-Verber results. To evaluate the effects of one synset’s frequency on another, we used Kendall’s tau (τ) rank order statistic (Kendall 1970). For example, was it the case that verbs under the synset *act* tend not to occur with verbs under the synset *think*? If so, do articles with this property fit a particular profile? In our results, we have information about synset frequency, where each of the 1236 articles in the corpus constitutes a sample. Table 3 shows the results of calculating Kendall’s τ with considerations for ranking ties, for all $\binom{10}{2} = 45$ pairing combinations of the top 10 most frequently occurring synsets. Correlations can range from -1.0 reflecting inverse correlation, to $+1.0$ showing direct correlation, i.e. the presence of one class increases as the presence of the correlated verb class increases. A τ value of 0 would show that the two variables’ values are independent of each other.

Results show a significant positive correlation between the synsets. The range of correlation is from .850 between the **communication** verb synset (*give, get, inform, ...*) and the **act** verb synset (*have, relate, give, ...*) to .238 between the **think** verb synset (*plan, study, give, ...*) and the **change state** verb synset (*fall, come, close, ...*).

These correlations show that frequent synsets do not behave independently of each other and thus confirm that the WordNet results are not an effective way to achieve document discrimination. Although the WordNet results were not discriminatory, we were still convinced that our initial hypothesis on the role of verbs in determining event profile was worth pursuing. We believe that these results are a by-product of lexical ambiguity and of the richness of the WordNet hierarchy. We thus decided to pursue a new approach to test our hypothesis, one which turned out to provide us with clearer and more robust results.

	act	com	chng	alter	infr	exps	thnk	judg	trnf
state	.407	.296	.672	.461	.286	.269	.238	.355	.268
trnsf	.437	.436	.251	.436	.251	.404	.369	.359	
judge	.444	.414	.435	.450	.340	.348	.427		
exps	.444	.414	.435	.397	.322	.432			
think	.444	.414	.435	.397	.398				
infr	.614	.649	.341	.380					
alter	.501	.454	.619						
chng	.496	.393							
comun	.850								

Table 3: Kendall's τ for frequent WordNet synsets.

Utilizing EVCA. A different approach to test the hypothesis was to use another semantic categorization method; we chose the semantic classes of Levin's EVCA as a basis for our next analysis.³ Levin's seminal work is based on the time-honored observation that verbs which participate in similar syntactic alternations tend to share semantic properties. Thus, the behavior of a verb with respect to the expression and interpretation of its arguments can be said to be, in large part, determined by its meaning. Levin has meticulously set out a list of syntactic tests (about 100 in all), which predict membership in no less than 48 classes, each of which is divided into numerous sub-classes. The rigor and thoroughness of Levin's study permitted us to encode our algorithm, **EVCA-Verber**, on a sub-set

³Strictly speaking, our classification is based on EVCA. Although many of our classes are precisely defined in terms of EVCA tests, we did impose some extensions. For example, support verbs are not an EVCA category.

of the EVCA classes, ones which were frequent in our corpus. First, we manually categorized the 100 most frequent verbs, as well as 50 additional verbs, which covers 56% of the verbs by token in the corpus. We subjected each verb to a set of strict linguistic tests, as shown in Table 4 and verified primary verb usage against the corpus.

Verb Class (sample verbs)	Sample Test
Communication (add, say, announce, ...)	(1) Does this involve a transfer of ideas? (2) X verbed "something."
Motion (rise, fall, decline, ...)	(1) *"X verbed without moving".
Agreement (agree, accept, occur, ...)	(1) "They verbed to join forces." (2) involves more than one participant.
Argument (argue, debate, ...)	(1) "They verbed (over) the issue." (2) indicates conflicting views. (3) involves more than one participant.
Causative (cause)	(1) X verbed Y (to happen/happened). (2) X brings about a change in Y.

Table 4: EVCA verb class test

Results from EVCA-Verber. In order to be able to compare article types and emphasize their differences, we selected articles that had the highest percentage of a particular verb class from each of the ten verb classes; we chose five articles from each EVCA class, yielding a total of 50 articles for analysis from the full set of 1236 articles. We observed that each class discriminated between different article types as shown in Table 5. In contrast to Table 2, the article types are well discriminated by verb class. For example, a concentration of **communication** class verbs (*say, report, announce, ...*) indicated that the article type was a general announcement of short or medium length, or a longer feature article with many opinions in the text. Articles high in **motion** verbs were also announcements, but differed from the communication ones, in that they were commonly postings of company earnings reaching a new high or dropping from last quarter. **Agreement** and **argument** verbs appeared in many of the same articles, involving issues of some controversy. However, we noted that articles with agreement verbs were a superset of the argument ones in that, in our corpus, argument verbs did not appear in articles concerning joint ventures and mergers. Articles marked by **causative** class verbs tended to be a bit longer, possibly reflecting prose on both the cause and effect of

a particular action. We also used **EVCA-Verber** to investigate articles marked by the absence of members of each verb class, such as articles lacking any verbs in the motion verb class. However, we found that absence of a verb class was not discriminatory.

Verb Class (sample verbs)	Article types (listed by frequency)
Communication (add, say, announce, ...)	issues, reports, opinions, editorials
Motion (rise, fall, decline, ...)	posted earnings, announcements
Agreement (agree, accept, concur, ...)	mergers, legal cases, transactions (without buying and selling)
Argument (argue, indicate, contend, ...)	legal cases, opinions
Causative (cause)	opinions, feature, editorials

Table 5: EVCA-based verb class results.

Evaluation of EVCA verb classes. To strengthen the observations that articles dominated by verbs of one class reflect distinct article types, we verified that the verb classes behaved independently of each other. Correlations for EVCA classes are shown in Table 6. These show a markedly lower level of correlation between verb classes than the results for WordNet synsets, the range being from .265 between motion and aspectual verbs to -.026 for motion verbs and agreement verbs. These low values of τ for pairs of verb classes reflects the independence of the classes. For example, the **communication** and **experience** verb classes are weakly correlated; this, we surmise, may be due to the different ways opinions can be expressed, i.e. as factual quotes using **communication** class verbs or as beliefs using **experience** class verbs.

	comun	motion	agree	argue	exp	aspect	cause
appear	.122	.076	.077	.072	.182	.112	.037
cause	.093	.083	.000	.000	.073	.096	
aspect	.246	.265	.034	.110	.189		
exp	.260	.130	.054	.054			
argue	.162	.045	.033				
agree	.071	-.026					
motion	.259						

Table 6: Kendall's τ for EVCA based verb classes.

4 Results and Future Work.

Basis for WordNet and EVCA comparison. This paper reports results from two approaches, one using WordNet and other based

on EVCA classes. However, the basis for comparison must be made explicit. In the case of WordNet, all verb tokens ($n = 10K$) were considered in all senses, whereas in the case of EVCA, a subset of less ambiguous verbs were manually selected. As reported above, we covered 56% of the verbs by token. Indeed, when we attempted to add more verbs to EVCA categories, at the 59% mark we reached a point of difficulty in adding new verbs due to ambiguity, e.g. verbs such as *get*. Thus, although our results using EVCA are revealing in important ways, it must be emphasized that the comparison has some imbalance which puts WordNet in an unnaturally negative light. In order to accurately compare the two approaches, we would need to process either the same less ambiguous verb subset with WordNet, or the full set of all verbs in all senses with EVCA. Although the results reported in this paper permitted the validation of our hypothesis, unless a fair comparison between resources is performed, conclusions about WordNet as a resource versus EVCA class distinctions should not be inferred.

Verb Patterns. In addition to considering verb type frequencies in texts, we have observed that verb distribution and patterns might also reveal subtle information in text. Verb class distribution within the document and within particular sub-sections also carry meaning. For example, we have observed that when sentences with movement verbs such as *rise* or *fall* are followed by sentences with *cause* and then a telic aspectual verb such as *reach*, this indicates that a value rose to a certain point due to the actions of some entity. Identification of such sequences will enable us to assign functions to particular sections of contiguous text in an article, in much the same way that text segmentation program seeks identify topics from distributional vocabulary (Hearst, 1994; Kan et al., 1998). We can also use specific sequences of verbs to help in determining methods for performing semantic aggregation of individual clauses in text generation for summarization.

Future Work. Our plans are to extend the current research in terms of verb coverage and in terms of article coverage. For verbs, we plan to (1) increase the verbs that we cover to include phrasal verbs; (2) increase coverage of verbs by categorizing additional high frequency verbs into EVCA classes; (3) examine the effects of

increased coverage on determining article type. For articles, we plan to explore a general parser so we can test our hypothesis on additional texts and examine how our conclusions scale up. Finally, we would like to combine our techniques with other indicators to form a more robust system, such as that envisioned in Biber (1989) or suggested in Kessler et al. (1997).

Conclusion. We have outlined a novel approach to document analysis for news articles which permits discrimination of the event profile of news articles. The goal of this research is to determine the role of verbs in document analysis, keeping in mind that event profile is one of many factors in determining text type. Our results show that Levin's EVCA verb classes provide reliable indicators of article type within the news domain. We have applied the algorithm to WSJ data and have discriminated articles with five EVCA semantic classes into categories such as features, opinions, and announcements. This approach to document type classification using verbs has not been explored previously in the literature. Our results on verb analysis coupled with what is already known about NP identification convinces us that future combinations of information will be even more successful in categorization of documents. Results such as these are useful in applications such as passage retrieval, summarization, and information extraction.

References

- D. Appelt, J. Hobbs, J. Bear, D. Isreal, and M. Tyson. 1993. Fastus: A finite state processor for information extraction from real world text. In *Proceedings of the 13th International Joint Conference on Artificial Intelligence (IJCAI)*, Chambery, France.
- Regina Barzilay and Michael Elhadad. 1997. Using lexical chains for text summarization. In *Proceedings of the Intelligent Scalable Text Summarization Workshop (ISTS'97)*, ACL, Madrid, Spain.
- Douglas Biber. 1989. A typology of english texts. *Language*, 27:3-43.
- Christiane Fellbaum. 1990. English verbs as a semantic net. *International Journal of Lexicography*, 3(4):278-301.
- Maarti A. Hearst. 1994. Multi-paragraph segmentation of expository text. In *Proceedings of the 32th Annual Meeting of the Association of Computational Linguistics*.
- Evan Hill and John J. Breen. 1977. *Reporting & Writing the News*. Little, Brown and Company, Boston, Massachusetts.
- Graeme Hirst and David St-Onge. 1998. Lexical chains as representations of context for the detection and correction of malapropisms. *WordNet: An electronic lexical database and some of its applications*.
- Ray Jackendoff. 1983. *Semantics and Cognition*. MIT University Press, Cambridge, Massachusetts.
- Min-Yen Kan, Judith L. Klavans, and Kathleen R. McKeeown. 1998. Linear segmentation and segment relevance. Unpublished Manuscript.
- Jussi Karlgren and Douglass Cutting. 1994. Recognizing text genres with simple metrics using discriminant analysis. In *Fifteenth International Conference on Computational Linguistics (COLING '94)*, Kyoto, Japan.
- Maurice G. Kendall. 1970. *Rank Correlation Methods*. Griffin, London, England, 4th edition.
- Brent Kessler, Geoffrey Nunberg, and Hinrich Schütze. 1997. Automatic detection of text genre. In *Proceedings of the 35th Annual Meeting of the Association of Computational Linguistics*, Madrid, Spain.
- Beth Levin. 1993. *English Verb Classes and Alternations*. University of Chicago Press, Chicago, Ohio.
- Chin-Yew Lin and Eduard Hovy. 1997. Identifying topics by position. In *Proceedings of the 5th ACL Conference on Applied Natural Language Processing*, pages 283-290, Washington, D.C., April.
- Dekang Lin. 1993. University of Manitoba: Description of the NUBA System as Used for MUC-5. In *Proceedings of the Fifth Conference on Message Understanding MUC-5*, pages 263-275, Baltimore, Maryland. ARPA.
- Mitch Marcus et al. 1994. *The Penn Treebank: Annotating Predicate Argument Structure*. ARPA Human Language Technology Workshop.
- 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 (special issue)*, 3(4):235-312.
- Jane Morris and Graeme Hirst. 1991. Lexical coherence computed by thesaural relations as an indicator of the structure of text. *Computational Linguistics*, 17(1):21-42.
1992. *Message Understanding Conference — MUC*.
- Günter Neumann, Rolf Backofen, Judith Baur, Marcus Becker, and Christian Braun. 1997. An information extraction core system for real world german text processing. In *Proceedings of the 5th ACL Conference on Applied Natural Language Processing*, pages 209-216, Washington, D.C., April.
- David D. Palmer and David S. Day. 1997. A statistical profile of the named entity task. In *Proceedings of the 5th ACL Conference on Applied Natural Language Processing*, pages 190-193, Washington, D.C., April.
- Philip Resnik. 1993. *Selection and Information: A Class-Based Approach to Lexical Relationships*. Ph.D. thesis, Department of Computer and Information Science, University of Pennsylvania.
- Nina Wacholder, Yael Ravin, and Misook Choi. 1997. Disambiguation of proper names in text. In *Proceedings of the 5th ACL Conference on Applied Natural Language Processing*, volume 1, pages 202-209, Washington, D.C., April.