

Using the Web to Obtain Frequencies for Unseen Bigrams

Frank Keller*
University of Edinburgh

Mirella Lapata†
University of Sheffield

This paper shows that the web can be employed to obtain frequencies for bigrams that are unseen in a given corpus. We describe a method for retrieving counts for adjective-noun, noun-noun, and verb-object bigrams from the web by querying a search engine. We evaluate this method by demonstrating: (a) a high correlation between web frequencies and corpus frequencies; (b) a reliable correlation between web frequencies and plausibility judgments; (c) a reliable correlation between web frequencies and frequencies recreated using class-based smoothing; (d) a good performance of web frequencies in a pseudo-disambiguation task.

Introduction

In two recent papers, Banko and Brill (2001a; 2001b) criticize the fact that current NLP algorithms are typically optimized, tested, and compared on fairly small data sets (corpora with millions of words), even though data sets several orders of magnitude larger are available, at least for some NLP tasks. Banko and Brill (2001a; 2001b) experiment with context-sensitive spelling correction, a task for which large amounts of data can be obtained straightforwardly, as no manual annotation is required. They demonstrate that the learning algorithms typically used for spelling correction benefit significantly from larger training sets, and that their performance shows no sign of reaching an asymptote as the size of the training set increases.

Arguably, the largest data set that is available for NLP is the web,¹ which currently consists of at least 968 million pages.² Data retrieved from the web therefore provides enormous potential for training NLP algorithms, if Banko and Brill's (2001a; 2001b) findings for spelling corrections generalize; potential applications include tasks that involve word n -grams and simple surface syntax. There is a small body of existing research that tries to harness the potential of the web for NLP. Grefenstette and Nioche (2000) and Jones and Ghani (2000) use the web to generate corpora for languages where electronic resources are scarce, while Resnik (1999) describes a method for mining the web in order to obtain bilingual texts. Mihalcea and Moldovan (1999) and Agirre and Martinez (2000) use the web for word sense disambiguation, Volk (2001) proposes a method for resolving PP attachment ambiguities based on web data, Markert, Nissim, and Modjeska (2003) use the web for the resolution of nominal anaphora, and Zhu and Rosenfeld (2001) use

* School of Informatics, 2 Buccleuch Place, Edinburgh EH8 9LW, UK. E-mail: keller@inf.ed.ac.uk

† Department of Computer Science, 211 Portobello Street, Sheffield S1 4DP, UK. E-mail: mlap@dcs.shef.ac.uk

1 A reviewer points out that information providers such as Lexis Nexis (<http://www.lexisnexis.com/>) might have databases that are even larger than the web. Lexis Nexis provides full-text access to news sources (including newspapers, wire services, broadcast transcripts) and legal data (including case law, codes, regulations, legal news, law reviews).

2 This is the number of pages indexed by Google in March 2002, as estimated by Search Engine Showdown (see <http://www.searchengineshowdown.com/>).

web-based n -gram counts to improve language modeling.

A particularly interesting application is proposed by Grefenstette (1998), who uses the web for example-based machine translation. His task is to translate compounds from French into English, with corpus evidence serving as a filter for candidate translations. An example is the French compound *groupe de travail*. There are five translations of *groupe* and three translations for *travail* (in the dictionary that Grefenstette (1998) is using), resulting in 15 possible candidate translations. Only one of them, viz., *work group* has a high corpus frequency, which makes it likely that this is the correct translation into English. Grefenstette (1998) observes that this approach suffers from an acute data sparseness problem if the counts are obtained from a conventional corpus. However, as Grefenstette (1998) demonstrates, this problem can be overcome by obtaining counts through web searches, instead of relying on a corpus. Grefenstette (1998) therefore effectively uses the web as a way of obtaining counts for compounds that are sparse in a given corpus.

While this is an important initial result, it raises the question of the generality of the proposed approach to overcoming data sparseness. It remains to be shown that web counts are generally useful for approximating data that is sparse or unseen in a given corpus. It seems possible, for instance, that Grefenstette's (1998) results are limited to his particular task (filtering potential translations) or to his particular linguistic phenomenon (noun-noun compounds). Another potential problem is the fact that web counts are far more noisy than counts obtained from a well-edited, carefully balanced corpus. The effect of this noise on the usefulness of the web counts is largely unexplored.

Zhu and Rosenfeld (2001) use web-based n -gram counts for language modeling. They obtain a standard language model from a 103 million word corpus and employ web-based counts to interpolate unreliable trigram estimates. They compare their interpolated model against a baseline trigram language model (without interpolation) and show that the interpolated model yields an absolute reduction in word error rate of .93% over the baseline. Zhu and Rosenfeld's (2001) results demonstrate that the web can be a source of data for language modeling. It is not clear, however, whether their result carries over to tasks that employ linguistically meaningful word sequences (e.g., head-modifier pairs or predicate-argument tuples) rather than simply adjacent words. Furthermore, Zhu and Rosenfeld (2001) do not undertake any studies that evaluate web frequencies directly (i.e., without a task such as language modeling). This could be done, for instance, by comparing web frequencies to corpus frequencies, or to frequencies recreated by smoothing techniques.

The aim of the present paper is to generalize Grefenstette's (1998) and Zhu and Rosenfeld's (2001) findings by testing the hypothesis that the web can be employed to obtain frequencies for bigrams that are unseen in a given corpus. Instead of having a particular task in mind (which would introduce a sampling bias), we rely on sets of bigrams that are randomly selected from a corpus. We use a web-based approach for bigrams that encode meaningful syntactic relations, and obtain web frequencies not only for noun-noun bigrams, but also for adjective-noun and verb-object bigrams. We thus explore whether this approach generalizes to different predicate-argument combinations. We evaluate our web counts in four ways: (a) comparison with actual corpus frequencies from two different corpora, (b) comparison with human plausibility judgments, (c) comparison with frequencies recreated using class-based smoothing, and (d) performance in a pseudo-disambiguation task on data sets from the literature.

1 Obtaining Frequencies from the Web

1.1 Sampling Bigrams from the BNC

The data sets used in the present experiment were obtained from the British National Corpus (BNC; see Burnard (1995)). The BNC is a large, synchronic corpus, consisting of 90 million words of text and 10 million words of speech. The BNC is a balanced corpus, i.e., it was compiled so as to represent a wide range of present day British English. The written part includes samples from newspapers, magazines, books (both academic and fiction), letters, and school and university essays, among other kinds of text. The spoken part consists of spontaneous conversations, recorded from volunteers balanced by age, region, and social class. Other samples of spoken language are also included, ranging from business or government meetings to radio shows and phone-ins. The corpus represents many different styles and varieties, and is not limited to any particular subject field, genre, or register.

For the present study, the BNC was used to extract data for three types of predicate-argument relations. The first type is adjective-noun bigrams, where we assume that the noun is the predicate that takes the adjective as its argument.³ The second predicate-argument type we investigate is noun-noun compounds. Here we assume that the rightmost noun is the predicate that selects the leftmost noun as its argument (as compound nouns are generally right-headed in English). Thirdly, we included verb-object bigrams, where the verb is the predicate that selects the object as its argument. We only considered direct NP objects; the bigram consists of the verb and the head noun of the object. For each of the three predicate-argument relations, we gathered two data sets, one containing seen bigrams, i.e., bigrams that occur in the BNC, and one with unseen bigrams, i.e., bigrams that fail to occur in the BNC.

For the seen adjective-noun bigrams, we used the data of Lapata, McDonald, and Keller (1999), who compiled a set of 90 bigrams as follows. First, 30 adjectives were randomly chosen from a part-of-speech tagged and lemmatized version of the BNC so that each adjective had exactly two senses according to WordNet (Miller et al., 1990) and was unambiguously tagged as “adjective” 98.6% of the time. They used the part-of-speech tagged version that is made available with the BNC and was tagged using CLAWS4 (Leech, Garside, and Bryant, 1994), a probabilistic part-of-speech tagger, with error rate ranging from 3% to 4%. The lemmatized version of the corpus was obtained using Karp et al.’s (1992) morphological analyzer.

The 30 adjectives ranged in BNC frequency from 1.9 to 49.1 per million words, i.e., they covered the whole range from fairly infrequent to highly frequent items. Gsearch (Corley et al., 2001), a chart parser which detects syntactic patterns in a tagged corpus by exploiting a user-specified context free grammar and a syntactic query, was used to extract all nouns occurring in a head-modifier relationship with one of the 30 adjectives. Examples of the syntactic patterns the parser identified are given in Table 1. In the case of adjectives modifying compound nouns, only sequences of two nouns were included and the rightmost occurring noun was considered as the head. Bigrams involving proper nouns or low-frequency nouns (less than 10 per million words) were discarded. This was necessary as the bigrams were used in experiments involving native speakers (see Section 2.2), and we wanted to reduce the risk of including words unfamiliar to the experimental subjects. For each adjective, the set of bigrams was divided into three frequency bands based on an equal division of the range of log-transformed co-occurrence frequencies. Then one

³ This assumption is disputed in the theoretical literature. For instance, Pollard and Sag (1994) present an analysis where there is mutual selection between the noun and the adjective.

Table 1

Example of patterns used for the extraction of adjective-noun bigrams

Pattern	Example
A N	educational material
A Adv N	usual weekly classes
A N N	environmental health officers

bigram was chosen at random from each band. This procedure makes sure that the whole range of frequencies is represented in our sample.

Lapata, Keller, and McDonald (2001) compiled a set of 90 unseen adjective-noun bigrams using the same 30 adjectives. For each adjective, Gsearch was used to compile a list of all nouns that failed to co-occur in a head-modifier relationship with the adjective. Again, proper nouns and low-frequency nouns were discarded from this list. Then each adjective was paired with three randomly chosen nouns from its list of non-co-occurring nouns. Examples of seen and unseen adjective-noun bigrams are shown in Table 2.

For the present study, we applied the procedure used by Lapata, McDonald, and Keller (1999) and Lapata, Keller, and McDonald (2001) to noun-noun bigrams and to verb-object bigrams, creating a set of 90 seen and 90 unseen bigrams for each type of predicate-argument relationship. More specifically, 30 nouns and 30 verbs were chosen according to the same criteria proposed for the adjective study (i.e., minimal sense ambiguity and unambiguous part-of-speech). All nouns modifying one of the 30 nouns were extracted from the BNC using a heuristic which looks for consecutive pairs of nouns that are neither preceded nor succeeded by another noun (Lauer, 1995). His heuristic (see (1)) effectively avoids identifying as two word compounds noun sequences which are part of a larger compound.

$$(1) \quad C = \{(w_2, w_3) \mid w_1 w_2 w_3 w_4; w_1, w_4 \notin N; w_2, w_3 \in N\}$$

Here, $w_1 w_2 w_3 w_4$ denotes the occurrence of a sequence of four words and N is the set of words tagged as nouns in the corpus. C is the set of compounds identified by Lauer's (1995) heuristic.

Verb-object bigrams for the 30 preselected verbs were obtained from the BNC using Cass (Abney, 1996), a robust chunk parser designed for the shallow analysis of noisy text. The parser recognizes chunks and simplex clauses (i.e., sequences of non-recursive clauses) using a regular expression grammar and a part-of-speech tagged corpus, without attempting to resolve attachment ambiguities. It comes with a large-scale grammar for English and a built-in tool that extracts predicate-argument tuples out of the parse trees that Cass produces.

The parser's output was post-processed to remove bracketing errors and errors in identifying chunk categories that could potentially result in bigrams whose members do not stand in a verb-argument relationship. Tuples containing verbs or nouns attested in a verb-argument relationship only once were eliminated. Particle verbs were retained only if the particle was adjacent to the verb (e.g., *come off heroin*). Verbs followed by the preposition *by* and a head noun were considered instances of verb-subject relations. It was assumed that PPs adjacent to the verb headed by either of the prepositions *in*, *to*, *for*, *with*, *on*, *at*, *from*, *of*, *into*, *through*, and *upon* were prepositional objects (see Lapata (2001) for details on the filtering process). Only nominal heads were retained from the objects returned by the parser. As in the adjective study, noun-noun bigrams and verb-object bigrams with proper nouns or low-frequency nouns (less than 10 per million

Table 2

Example stimuli for seen and unseen adjective-noun, noun-noun, and verb-object bigrams (with log-transformed BNC counts)

adjective-noun bigrams							
adjective	high		medium		low		unseen
hungry	animal	1.79	pleasure	1.38	application	0	tradition, innovation, prey
guilty	verdict	3.91	secret	2.56	cat	0	system, wisdom, wartime
naughty	girl	2.94	dog	1.6	lunch	.69	regime, rival, protocol
noun-noun bigrams							
high		medium		low		unseen	head noun
process	1.14	user	.95	gala	0	collection, clause, coat	directory
television	1.53	satellite	.95	edition	0	chain, care, vote	broadcast
plasma	1.78	nylon	1.20	unit	.60	fund, theology, minute	membrane
verb-object bigrams							
verb	high		medium		low		unseen
fulfill	obligation	3.87	goal	2.20	scripture	.69	participant, muscle, grade
intensify	problem	1.79	effect	1.10	alarm	0	score, quota, chest
choose	name	3.74	law	1.61	series	1.10	lift, bride, listener

words) were discarded. The sets of noun-noun and verb-object bigrams were divided into three frequency bands and one bigram was chosen at random from each band.

The procedure described by Lapata, Keller, and McDonald (2001) was followed for creating sets of unseen noun-noun and verb-object bigrams: for each of noun or verb, we compiled a list of all nouns with which it failed to co-occur within a noun-noun or verb-object bigram in the BNC. Again, Lauer’s (1995) heuristic and Abney’s (1996) partial parser were used to identify bigrams, and proper nouns and low-frequency nouns were excluded. For each noun and verb, three bigrams were randomly selected from the set of their non-co-occurring nouns. Table 2 lists examples for the seen and unseen noun-noun and verb-object bigrams generated by this procedure.

The extracted bigrams are in several respects an imperfect source of information about adjective-noun or noun-noun modification and verb-object relations. First notice that both Gsearch and Cass detect syntactic patterns on part-of-speech tagged corpora. This means that parsing errors are likely to result because of tagging mistakes. Second, even if one assumes perfect tagging, the heuristic nature of our extraction procedures may introduce additional noise or miss bigrams for which detailed structural information would be needed.

For instance, our method for extracting adjective-noun pairs ignores cases where the adjective modifies noun sequences of length larger than two. The heuristic in (1) considers only two word noun sequences. Abney’s (1996) chunker recognizes basic syntactic units without resolving attachment ambiguities or recovering missing information (such as traces resulting from the movement of constituents). Although parsing is robust and fast (since unlike traditional parsers no global optimization takes place), the identified verb-argument relations are undoubtedly somewhat noisy given the errors inherent in the part-of-speech tagging and chunk recognition procedure. When evaluated against manually annotated data, Abney’s (1996) parser identified chunks with 87.9% precision and 87.1% recall. The parser further achieved a per-word accuracy of 92.1% (where per-word accuracy includes the chunk category and chunk length identified correctly).

Despite their imperfect output, heuristic methods for the extraction of syntactic relations are relatively common in statistical NLP. Several statistical models employ frequencies obtained from the output of partial parsers and other heuristic methods;

Table 3Examples of dependencies generated by MINIPAR for *The fat cat ate the door mat*

Head	Relation	Modifier	Description
cat	N:det:Det	the	determiner of noun
cat	N:mod:A	fat	adjective modifier of noun
eat	V:subj:N	cat	subject of verb
eat	V:obj:N	mat	object of verb
mat	N:det:Det	the	determiner of noun
mat	N:nn:N	door	prenominal modifier of noun

these include models for disambiguating the attachment site of prepositional phrases (Hindle and Rooth, 1993; Ratnaparkhi, 1998), models for interpreting compound nouns (Lauer, 1995; Lapata, 2002) and polysemous adjectives (Lapata, 2001), models for the induction of selectional preferences (Abney and Light, 1999), methods for automatically clustering words according to their distribution in particular syntactic contexts (Pereira, Tishby, and Lee, 1993), automatic thesaurus extraction (Grefenstette, 1994; Curran, 2002), and similarity-based models of word co-occurrence probabilities (Lee, 1999; Dagan, Lee, and Pereira, 1999). In this paper we investigate alternative ways for obtaining bigram frequencies that are potentially useful for such models despite the fact that some of these bigrams are identified in a heuristic manner and may be noisy.

1.2 Sampling Bigrams from the NANTC

We also obtained corpus counts from a second corpus, the North American News Text Corpus (NANTC). This corpus differs in several important respects from the BNC. It is substantially larger, as it contains 350 million words of text. Also, it is not a balanced corpus, as it only contains material from one genre, viz., news text. However, the text originates from a variety of sources (Los Angeles Times, Washington Post, New York Times News Syndicate, Reuters News Service, and Wall Street Journal). While the BNC covers British English, the NANTC covers American English. All these differences mean that the NANTC provides a second, independent standard against which to compare web counts. At the same time the correlation found between the counts obtained from the two corpora can serve as an upper limit for the correlation that we can expect between corpus counts and web counts.

The NANTC corpus was parsed using MINIPAR (Lin, 1994; Lin, 2001) a broad coverage parser for English. MINIPAR employs a manually constructed grammar and a lexicon derived from WordNet with the addition of proper names (130,000 entries in total). Lexicon entries contain part-of-speech and subcategorization information. The grammar is represented as a network of 35 nodes (i.e., grammatical categories) and 59 edges (i.e., types of syntactic (dependency) relationships). MINIPAR employs a distributed chart parsing algorithm. Instead of a single chart, each node in the grammar network maintains a chart containing partially built structures belonging to the grammatical category represented by the node. Grammar rules are implemented as constraints associated with the nodes and edges.

The output of MINIPAR is a dependency tree which represents the dependency relations between words in a sentence. Table 3 shows a subset of the dependencies MINIPAR outputs for the sentence *The fat cat ate the door mat*. In contrast to Gsearch and Cass, MINIPAR produces all possible parses for a given sentence. The parses are ranked according to the product of the probabilities of their edges and the most likely parse is returned.

Table 4

Log-transformed NANTC counts for seen adjective-noun, noun-noun and verb-object bigrams

adjective-noun bigrams						
adjective	high	medium	low			
hungry	animal	.90	pleasure	-.30	application	.60
guilty	verdict	2.82	secret	.95	cat	-.30
naughty	girl	.69	dog	-.30	lunch	-.30
noun-noun bigrams						
high	medium	low	head noun			
process	-.30	user	-.30	gala	-.30	directory
television	2.70	satellite	-.30	edition	-.30	broadcast
plasma	-.30	nylon	0	unit	0	membrane
verb-object bigrams						
verb	high	medium	low			
fulfill	obligation	2.38	goal	2.04	scripture	-.30
intensify	problem	1.20	effect	.60	alarm	-.30
choose	name	2.25	law	.90	series	.48

Lin (1998) evaluated the parser on the SUSANNE corpus (Sampson, 1995), a domain independent corpus of British English, and achieved a recall of 79% and precision of 89% on the dependency relations.

For our experiments, we concentrated solely on adjective-noun, noun-noun, and verb object relations (denoted as $N:mod:A$, $N:nn:N$, and $V:obj:N$ in Table 3). From the syntactic analysis provided by the parser, we extracted all occurrences of bigrams that were attested both in the BNC and the NANTC corpus. This way, we obtained NANTC frequency counts for the bigrams that we had randomly selected from the BNC. Table 4 shows the NANTC counts for the set of seen bigrams from Table 2.

Due to the differences in the extraction methodology (chunking versus full parsing) and the text genre (balanced corpus versus news text), we expect that some BNC bigrams will not be attested in the NANTC corpus. More precisely, zero frequencies were returned for 23 adjective-noun, 16 verb-noun, and 37 noun-noun bigrams. The fact that more zero frequencies were observed for noun-noun bigrams is perhaps not surprising considering the ease with which novel compounds are created (Levi, 1978). We adjusted the zero counts by setting them to .5. This was necessary as all further analyses were carried out on log-transformed frequencies (see Table 4).

1.3 Obtaining Web Counts

Web counts for bigrams were obtained using a simple heuristic based on queries to the search engines Altavista and Google. All search terms took into account the inflectional morphology of nouns and verbs.

The search terms for verb-object bigrams matched not only cases in which the object was directly adjacent to the verb (e.g., *fulfill obligation*), but also cases where there was an intervening determiner (e.g., *fulfill the/an obligation*). The following search terms were used for adjective-noun, noun-noun, and verb-object bigrams, respectively:

- (2) "A N", where A is the adjective and N is the singular or plural form of the noun.
- (3) "N1 N2" where N1 is the singular form of the first noun and N2 is the singular or plural form of the second noun.

Table 5

Number of zero counts returned by the queries to search engines and in the NANTC (for bigrams unseen in the BNC)

	adj-noun	noun-noun	verb-object
Altavista	2	9	1
Google	2	5	0
NANTC	76	82	78

- (4) "V Det N" where V is the infinitive, singular present, plural present, past, perfect, or gerund for of the verb, Det is the determiner *the*, *a* or the empty string, and N is the singular or plural form of the noun.

Note that all searches were for exact matches, which means that the words in the search terms had to be directly adjacent to score a match. This is encoded using quotation marks to enclose the search term. All our search terms were in lower case. We searched the whole web (as indexed by Altavista and Google), i.e., the queries were not restricted to pages in English.

Based on the web searches, we obtained bigram frequencies by adding up the number of pages that matched the morphologically expanded forms of the search terms (see the patterns in (2)–(4)). This process can be automated straightforwardly using a script that generates all the search terms for a given bigram, issues an Altavista or Google query for each of the search terms, and then adds up the resulting number of matches for each bigram. We applied this process to all the bigrams in our data set, covering seen and unseen adjective-noun, noun-noun, and verb-object bigrams, i.e., a set of 540 bigrams in total. The queries were carried out in January 2003 (and thus are higher than the counts reported in Keller, Lapata, and Ourioupina (2002), which were generated about a year earlier).

For some bigrams that were unseen in the BNC, our web-based procedure returned zero counts, i.e., there were no matches in the web searches. It is interesting to compare the web and NANTC with respect to zero counts: both data sources are larger than the BNC, and hence should be able to mitigate the data sparseness problem to a certain extent. Table 5 lists the number of zero counts for both web search engines and compares them to the number of bigrams that yielded no matches in the NANTC. We observe that the web counts are substantially less sparse than the NANTC counts: in the worst case there are nine bigrams for which our web queries returned no matches (10% of the data), while up to 82 bigrams were unseen in the NANTC (91% of the data). Recall that the NANTC is 3.5 times larger than the BNC, which does not seem to be enough to substantially mitigate data sparseness. All further analyses were carried out on log-transformed frequencies, hence we adjusted zero counts by setting them to .5.

Table 6 lists the descriptive statistics for the bigram counts we obtained using Altavista and Google. For comparison, this table also contains the descriptive statistics for the BNC and NANTC counts (for seen bigrams only), and the counts recreated using class-based smoothing (see Section 2.3 for details on the recreated frequencies).

From these data, we computed the average factor by which the web counts are larger than the BNC counts. The results are given in Table 7 and indicate that the Altavista counts are between 550 and 691 times larger than the BNC counts, while the Google counts are between 1064 and 1306 times larger than the BNC counts. As we know the size of the BNC (100 million words), we can use these figures to estimate the number of words available on the web: between 55.0 and 69.1 billion words for Altavista, and

Table 6

Descriptive statistics for web counts, corpus counts, and counts recreated using class-based smoothing (log-transformed)

	seen bigrams											
	adj-noun				noun-noun				verb-object			
	Min	Max	Mean	SD	Min	Max	Mean	SD	Min	Max	Mean	SD
Altavista	1.15	5.84	3.72	1.02	.60	6.16	3.52	1.22	.48	5.86	3.42	1.13
Google	1.54	6.11	4.01	1.01	.90	6.30	3.80	1.23	.60	5.96	3.70	1.11
BNC	0	2.19	.89	.69	0	2.14	.74	.64	0	2.55	.68	.58
NANTC	-.30	2.84	.84	.96	-.30	3.02	.56	.94	-.30	3.73	1.90	.98
Smoothing	-.06	2.32	1.28	.51	-.70	1.71	.30	.61	-.51	2.07	.53	.57

	unseen bigrams											
	adj-noun				noun-noun				verb-object			
	Min	Max	Mean	SD	Min	Max	Mean	SD	Min	Max	Mean	SD
Altavista	-.30	5.00	1.50	.99	-.30	3.97	1.20	1.14	-.30	3.88	1.55	1.06
Google	-.30	4.11	1.79	.95	-.30	4.15	1.60	1.12	0	4.19	1.90	1.04
Smoothing	-.03	2.10	1.25	.46	-1.01	1.93	.28	.66	-.70	1.95	.53	.58

Table 7

Average factor by which the web counts are larger than the BNC counts (seen bigrams)

	adj-noun	noun-noun	verb-object
Altavista	665	691	550
Google	1306	1151	1064

between 106.4 and 139.6 billion words for Google. These estimates are in the same order of magnitude as Grefenstette and Nioche’s (2000) estimate that 48.1 billion words of English are available on the web (based on Altavista counts compiled in February 2000). They also agree with Zhu and Rosenfeld’s (2001) estimate that the effective size of the web is between 79 and 108 billion words (based on Altavista, Lycos, and FAST counts; no date given).

1.4 Potential Sources of Noise in Web Counts

The method we used to retrieve web counts is based on very simple heuristics; it is thus inevitable that the counts generated will contain a certain amount of noise. In this section we discuss a number of potential sources of such noise.

An obvious limitation of our method is that it relies on the page counts returned by the search engines; we do not download the pages themselves for further processing. Note that many of the bigrams in our sample are very frequent (up to 10^6 matches, see the “Max” columns in Table 6), hence the effort involved in downloading all pages would be immense (though methods for downloading a representative could probably be devised).

Our approach estimates web frequencies based not on bigram counts directly, but on page counts. In other words, it ignores the fact that a bigram can occur more than once on a given web page. This approximation is justified, as Zhu and Rosenfeld (2001) demonstrated for unigrams, bigrams, and trigrams: page counts and n -gram counts are highly correlated on a log-log scale. This result is based on queries to Altavista, a search engine which returns both the number of pages and the overall number of matches for a

given query.⁴

Another important limitation of our approach is due to fact that both Google and Altavista remove punctuation and capitalization, even if the search term is quoted. This can lead to false positives, for instance if the match crosses a phrase boundary, such as in (5) which matches *hungry prey*. Other false positives can be generated by page titles and links, such as the examples (6) and (7) which match *edition broadcast*.⁵

- (5) The lion will kill only when it's **hungry**. **Prey** can usually sense when lions are hunting.
- (6) 10th **Edition Broadcast** Products Catalog (as a page title)
- (7) Issue/**Edition/Broadcast** (as a link)

The fact that our method does not download web pages means that no tagging, chunking, and parsing can be carried out to ensure that the matches are correct. Instead we rely on the simple adjacency of the search terms, which is enforced by using quoted queries (see Section 1.3 for details). This means that we miss any non-adjacent matches, even though a chunker or parser (such as the one used for extracting BNC or NANTC bigrams) would find them. An example is an adjective-noun bigram in which an adverbial intervenes between the adjective and the noun (see Table 1).

Furthermore, the absence of tagging, chunking, and parsing can also generate false positives, in particular for queries containing words with part-of-speech ambiguity. (Recall that our bigram selection procedure ensures that the predicate word, but not the argument word, is unambiguous in terms of its POS tagging in the BNC.) As an example consider *process directory*, which in our data set is a noun-noun bigram (see Table 2). One of the matches returned by Google is (8), where *process* is a verb. Another example is *fund membrane*, which is a noun-noun bigram in our data set, but which matches (9) in Google.

- (8) The global catalog server's function is to **process directory** searches for the entire forest.
- (9) Green grants **fund membrane** technology.

Another source of noise is the fact that Google (but not Altavista) will sometimes return pages that do not include the search term at all. This can happen if the search term is contained in a link to the page (but not in the page itself).

As we did not limit our web searches to English (even though many search engines now allow to set the target language for a search), there is also a risk that false positives are generated by crosslinguistic homonyms, i.e., by words of other languages that are spelled in the same way as the English words in our data sets. However, this problem is mitigated by the fact that English is by far the most common language on the web, as shown by Grefenstette and Nioche (2000). Also, the chance of two such homonyms forming a valid bigram in another language is probably fairly small.

To summarize, web counts are certainly less sparse than the counts in a corpus of a fixed size (see Section 1.3). However, web counts are also likely to be significantly more noisy than counts obtained from a carefully tagged and chunked or parsed corpus, as the examples in this section show. It is therefore essential to carry out a comprehensive evaluation of the web counts generated by our method. This is the topic of the next section.

⁴ Note that this feature of Altavista has since been discontinued; hence in the present paper we had no option but to use page counts. However, Keller, Lapata, and Ourioupina (2002) used Altavista match counts (instead of page counts) on the same data sets; their results agree with the ones reported in the present paper very closely.

⁵ Some of the examples in (5)–(9) were kindly provided by a reviewer.

Table 8

Correlation of BNC counts with web counts (seen bigrams)

	adj-noun	noun-noun	verb-object
Altavista	.847**	.720**	.762**
Google	.850**	.720**	.766**
Smoothing	.248*	.277**	.342**
* $p < .05$ (one-tailed)		** $p < .01$ (one-tailed)	

Table 9

Correlation of NANTC counts with web counts (seen bigrams)

	adj-noun	noun-noun	verb-object
Altavista	.712**	.667**	.788**
Google	.712**	.662**	.787**
BNC	.710**	.672**	.814**
Smoothing	.338**	.317**	.263*
* $p < .05$ (one-tailed)		** $p < .01$ (one-tailed)	

2 Evaluation

2.1 Evaluation Against Corpus Frequencies

Since web counts can be relatively noisy, as discussed in the previous section, it is crucial to determine if there is a reliable relationship between web counts and corpus counts. Once this is assured, we can explore the usefulness of web counts for overcoming data sparseness. We carried out a correlation analysis to determine if there is a linear relationship between BNC and NANTC counts and Altavista and Google counts. All correlation coefficients reported in this paper refer to Pearson’s r .⁶ All results were obtained on log-transformed counts.⁷

Table 8 lists the results of correlating web counts with corpus counts from the BNC, the corpus that our bigrams were sampled from (see Section 1.1). A high correlation coefficient was obtained across the board, ranging from .720 to .847 for Altavista counts and from .720 to .850 for Google counts. This indicates that web counts approximate BNC counts for the three types of bigrams under investigation. Note that there is almost no difference between the correlations achieved using Google and Altavista counts.

It is important to check that these results are also valid for counts obtained from other corpora. We therefore correlated our web counts with the counts obtained from NANTC, a corpus that is larger than the BNC, but is drawn from a single genre, viz., news text (see Section 1.2). The results are listed in Table 9. We find that Google and Altavista counts also correlate significantly with NANTC counts. The correlation coefficients range from .667 to .788 for Altavista and from .662 to .787 for Google. Again, there is virtually no

⁶ Correlation analysis is a way of measuring the degree of linear association between two variables.

Effectively, we are fitting a linear equation $y = ax + b$ to the data; this means that the two variables x and y (which in our case represent frequencies or judgments) can still differ by a multiplicative constant a and an additive constant b , even if they are highly correlated.

⁷ It is well-known that corpus frequencies have a Zipfian distribution. Log-transforming them is a way of normalizing the counts before applying statistical tests. We apply correlation analysis on the log-transformed data, which is equivalent to computing a log-linear regression coefficient on the untransformed data.

difference between the two search engines. We also observe that the correlation between web counts and BNC is generally slightly higher than the correlation between web counts and NANTC counts. We carried out one-tailed t -tests to determine if the differences in the correlation coefficients were significant. We found that both Altavista counts ($t(87) = 3.11$, $p < .01$) and Google counts ($t(87) = 3.21$, $p < .01$) were significantly better correlated with BNC counts than with NANTC counts for adjective-noun bigrams. The difference in correlation coefficients was not significant for noun-noun and verb-object bigrams, for either search engine.

Table 9 also lists the correlations between BNC counts and NANTC counts. The inter-corpus correlation can be regarded as an upper limit for the correlations we can expect between counts from two corpora that differ in size and genre, and that were obtained using different extraction methods. The correlation between Altavista and Google counts and NANTC counts reached the upper limit for all three bigram types (one-tailed t -tests found no significant differences between the correlation coefficients). The correlation between BNC counts and web counts reached the upper limit for noun-noun and verb-object bigrams (no significant differences for either search engine), and significantly exceeded it for adjective-noun bigrams for Altavista ($t(87) = 3.16$, $p < .01$) and Google ($t(87) = 3.26$, $p < .01$).

We conclude that simple heuristics (see Section 1.3) are sufficient to obtain useful frequencies from the web; it seems that the large amount of data available for web counts outweighs the associated problems (noisy, unbalanced, etc.). We found that web counts were highly correlated with frequencies from two different corpora. Furthermore, web counts and corpus counts are as highly correlated as counts from two different corpora (which can be regarded as an upper bound).

Note that Tables 8 and 9 also contain the correlation coefficients obtained when comparing corpus frequencies with frequencies that were recreated using class-based smoothing, using the BNC as a training corpus (after removing the seen bigrams). This will be discussed in more detail in Section 2.3.

2.2 Evaluation Against Plausibility Judgments

Previous work has demonstrated that corpus counts correlate with human plausibility judgments for adjective-noun bigrams. This result holds for both seen bigrams (Lapata, McDonald, and Keller, 1999) and for unseen bigrams whose counts were recreated using smoothing techniques (Lapata, Keller, and McDonald, 2001). Based on these findings, we decided to evaluate our web counts on the task of predicting plausibility ratings. If the web counts for bigrams correlate with plausibility judgments, then this indicates that the counts are valid, in the sense of being useful for predicting the intuitive plausibility of predicate-argument pairs. The degree of correlation between web counts and plausibility judgments is an indicator of the quality of the web counts (compared to corpus counts or counts recreated using smoothing techniques).

2.2.1 Method For seen and unseen adjective-noun bigrams, we used the two sets of plausibility judgments collected by Lapata, McDonald, and Keller (1999) and Lapata, Keller, and McDonald (2001), respectively. We conducted four additional experiments to collect judgments for noun-noun and verb-object bigrams, both seen and unseen. The experimental method was the same for all six experiments.

Materials. The experimental stimuli were based on the six sets of seen or unseen bigrams extracted from the BNC as described in Section 1.1 (adjective-noun, noun-noun, and verb-object bigrams). In the adjective-noun and noun-noun case, the stimuli simply consisted of the bigrams. In the verb-object case, the bigrams were embedded in a short

sentence to make them more natural: a proper name subject was added.

Procedure. The experimental paradigm was magnitude estimation (ME), a technique standardly used in psychophysics to measure judgments of sensory stimuli (Stevens, 1975), which Bard, Robertson, and Sorace (1996) and Cowart (1997) have applied to the elicitation of linguistic judgments. The ME procedure requires subjects to estimate the magnitude of physical stimuli by assigning numerical values proportional to the stimulus magnitude they perceive. In contrast to the 5- or 7-point scale conventionally used to measure human intuitions, ME employs an interval scale, and therefore produces data for which parametric inferential statistics are valid.

ME requires subjects to assign numbers to a series of linguistic stimuli in a proportional fashion. Subjects are first exposed to a modulus item, which they assign an arbitrary number. All other stimuli are rated proportional to the modulus. In this way, each subject can establish their own rating scale, thus yielding maximally fine-grained data and avoiding the known problems with the conventional ordinal scales for linguistic data (Bard, Robertson, and Sorace, 1996; Cowart, 1997; Schütze, 1996).

The experiments reported in this paper were carried out using the WebExp software package (Keller et al., 1998). A series of previous studies has shown that data obtained using WebExp closely replicates results obtained in a controlled laboratory setting; this was demonstrated for acceptability judgments (Keller and Alexopoulou, 2001), co-reference judgments (Keller and Asudeh, 2001), and sentence completions (Corley and Scheepers, 2002).

In the present experiments, subjects were presented with bigram pairs and were asked to rate the degree of plausibility proportional to a modulus item. They first saw a set of instructions that explained the ME technique and the judgment task. The concept of plausibility was not defined, but examples of plausible and implausible bigrams were given (different examples for each stimulus set). Then subjects had to fill in a questionnaire with basic demographic information. The experiment proper consisted of three phases: (a) a calibration phase designed to familiarize subjects with the task; here, they had to estimate the length of five horizontal lines; (b) a practice phase in which subjects judged the plausibility of eight bigrams (similar to the ones in the stimulus set); (c) the main experiment, in which each subject judged one of the six stimulus sets (90 bigrams). The stimuli were presented in random order, with a new randomization being generated for each subject.

Subjects. A separate experiment was run for each set of stimuli. The number of subjects per experiment is listed in Table 10 (in the column labeled “N”). All Subjects were self-reported native speakers of English; they were recruited by postings to newsgroups and mailing lists. Participation was voluntary and unpaid.

WebExp collects by-item response time data; subjects whose response times were very short or very long were excluded from the sample, as they are unlikely to have completed the experiment adequately. We excluded the data of subjects that had participated more than once in the same experiment, based on their demographic data and on their internet connection data, which is logged by WebExp.

2.2.2 Results and Discussion The experimental data were normalized by dividing each numerical judgment by the modulus value that the subject had assigned to the reference sentence. This operation creates a common scale for all subjects. Then the data were transformed by taking the decadic logarithm. This transformation ensures that the judgments are normally distributed and is standard practice for magnitude estimation data (Bard, Robertson, and Sorace, 1996; Cowart, 1997; Stevens, 1975). All

Table 10

Descriptive statistics for plausibility judgments (log-transformed); N is the number of subjects used in each experiment

	adj-noun bigrams					noun-noun bigrams					verb-object bigrams				
	N	Min	Max	Mean	SD	N	Min	Max	Mean	SD	N	Min	Max	Mean	SD
Seen	30	-.85	.11	-.13	.22	25	-.15	.69	.40	.21	27	-.52	.45	.12	.24
Unseen	41	-.56	.37	-.07	.20	25	-.49	.52	-.01	.23	21	-.51	.28	-.16	.22

further analyses were conducted on the normalized, log-transformed judgments.

Table 10 lists the descriptive statistics for all six judgment experiments: the original experiments by Lapata, McDonald, and Keller (1999) and Lapata, Keller, and McDonald (2001) for adjective-noun bigrams, and our new ones for noun-noun and verb-object bigrams.

We used correlation analysis to compare corpus counts and web counts with plausibility judgments. Table 11 (top half) lists the correlation coefficients that were obtained when correlating log-transformed web counts (Altavista and Google) and corpus counts (BNC and NANTC) with mean plausibility judgments for seen adjective-noun, noun-noun, and verb-object bigrams.

The results show that both Altavista and Google counts correlate well with plausibility judgments for seen bigrams. The correlation coefficient for Altavista ranges from .641 to .700; for Google, it ranges from .624 to .692. The correlations for the two search engines are very similar, which is also what we found in Section 2.1 for the correlations between web counts and corpus counts.

Note that the web counts consistently achieve a higher correlation with the judgments than the BNC counts, which range from .488 to .569. We carried out a series of one-tailed t -tests to determine if the differences between the correlation coefficients for the web counts and the correlation coefficients for the BNC counts were significant. For the adjective-noun bigrams, the Altavista coefficient was significantly higher than the BNC coefficient ($t(87) = 1.76, p < .05$), while the difference between the Google coefficient and the BNC coefficient failed to reach significance. For the noun-noun bigrams, both the Altavista and the Google coefficients were significantly higher than the BNC coefficient ($t(87) = 3.11, p < .01$ and $t(87) = 2.95, p < .01$). Also for the verb-object bigrams, both the Altavista coefficient and the Google coefficient were significantly higher than the BNC coefficient ($t(87) = 2.64, p < .01$ and $t(87) = 2.32, p < .05$).

A similar picture was found for the NANTC counts. Again, the web counts outperform the corpus counts in predicting plausibility. For the adjective-noun bigrams, both the Altavista and the Google coefficient were significantly higher than the NANTC coefficient ($t(87) = 1.97, p < .05$; $t(87) = 1.81, p < .05$). For the noun-noun bigram, the Altavista coefficient was higher than the NANTC coefficient ($t(87) = 1.64, p < .05$), but the Google coefficient was not significantly different from the NANTC coefficient. For verb-object bigrams, the difference was significant for both search engines ($t(87) = 2.74, p < .01$; $t(87) = 2.38, p < .01$).

In sum, for all three types of bigrams, the correlation coefficients achieved with Altavista were significantly higher than the ones achieved by both BNC and NANTC. Google counts outperformed corpus counts for all bigrams with the exception of adjective-noun counts from the BNC, and noun-noun counts from the NANTC.

Table 11 (bottom half) lists the correlations coefficients obtained by comparing log-transformed judgments with log-transformed web counts for unseen adjective-noun, noun-noun, and verb-object bigrams. We observe that the web counts consistently show

Table 11

Correlation of plausibility judgments with web counts, corpus counts, and counts recreated using class-based smoothing; Agreement is the inter-subject agreement on the judgment task

seen bigrams			
	adj-noun	noun-noun	verb-object
Altavista	.650**	.700**	.641**
Google	.641**	.692**	.624**
BNC	.569**	.517**	.488**
NANTC	.526**	.597**	.491**
Smoothing	.329**	.318**	.223*
Agreement	.630**	.641**	.604**
unseen bigrams			
Altavista	.480**	.578**	.551**
Google	.473**	.595**	.520**
Smoothing	.342**	.372**	.298**
Agreement	.550**	.570**	.640**
* $p < .05$ (one-tailed)		** $p < .01$ (one-tailed)	

a significant correlation with the judgments, the coefficient ranging from .480 to .578 for Altavista counts, and from .473 to .595 for the Google counts. Table 11 also lists the correlations between plausibility judgments and counts recreated using class-based smoothing, which we will discuss in Section 2.3.

An important question is how well humans agree when judging the plausibility of adjective-noun, noun-noun and verb-noun bigrams. Inter-subject agreement gives an upper bound for the task and allows us to interpret how well our web-based method does in relation to humans. To calculate inter-subject agreement we used leave one-out resampling. This technique is a special case of n -fold crossvalidation (Weiss and Kulikowski, 1991) and has been previously used for measuring how well humans agree on judging semantic similarity (Resnik, 1999; Resnik, 2000).

For each subject group, we divided the set of the subjects' responses with size n into a set of size $n - 1$ (i.e., the response data of all but one subject) and a set of size one (i.e., the response data of a single subject). We then correlated the mean ratings of the former set with the ratings of the latter. This was repeated n times (see the number of participants in Table 6); the mean of the correlation coefficients is shown in Table 11 under the heading "Agreement" for the seen and unseen bigrams.

For both seen and unseen bigrams, we found no significant difference between the upper bound (inter-subject agreement) and the correlation coefficients reached using either Altavista or Google counts. This holds for all three types of bigrams. The same picture emerged for the BNC and NANTC counts: these correlation coefficients were not significantly different from the upper limit, for all three types of bigrams, both for seen and for unseen bigrams.

To conclude, our evaluation demonstrated that web counts reliably predict human plausibility judgments, both for seen and for unseen predicate-argument bigrams. Altavista counts for seen bigrams are a better predictor of human judgments than BNC and NANTC counts. These results show that our heuristic method yields valid frequencies; the simplifications we made in obtaining the web counts (see Section 1.3), as well as the fact that web data are noisy (see Section 1.4), seem to be outweighed by the fact that the web is up to a thousand times larger than the BNC.

Table 12

Correlation of counts recreated using class-based smoothing with web counts

seen bigrams			
	adj-noun	noun-noun	verb-object
Altavista	.344**	.362**	.361**
Google	.330**	.343**	.349**
unseen bigrams			
Altavista	.439**	.386**	.412**
Google	.444**	.421**	.397**
* $p < .05$ (one-tailed)		** $p < .01$ (one-tailed)	

2.3 Evaluation Against Class-based Smoothing

The evaluation in the last two sections established that web counts are useful for approximating corpus counts and for predicting plausibility judgments. As a further step in our evaluation, we correlated web counts with counts recreated by applying a class-based smoothing method on the BNC.

We recreated co-occurrence frequencies for predicate-argument bigrams using a simplified version of Resnik's (1993) selectional association measure proposed by Lapata, Keller, and McDonald (2001). In a nutshell, this measure replaces Resnik's (1993) information-theoretic approach with a simpler measure which makes no assumptions with respect to the contribution of a semantic class to the total quantity of information provided by the predicate about the semantic classes of its argument. It simply substitutes the argument occurring in the predicate-argument bigram with the concept by which it is represented in the WordNet taxonomy. Predicate-argument co-occurrence frequency is estimated by counting the number of times the concept corresponding to the argument is observed to co-occur with the predicate in the corpus. Because a given word is not always represented by a single class in the taxonomy (i.e., the argument co-occurring with a predicate can generally be the realization of one of several conceptual classes), Lapata, Keller, and McDonald (2001) constructed the frequency counts for a predicate-argument bigram for each conceptual class by dividing the contribution from the argument by the number of classes to which it belongs. They demonstrate that the counts recreated using this smoothing technique correlate significantly with plausibility judgments for adjective-noun bigrams. They also show that this class-based approach outperforms distance-weighted averaging (Dagan, Lee, and Pereira, 1999), a smoothing method that recreates unseen word co-occurrences on the basis of distributional similarity (without relying on a predefined taxonomy), in predicting plausibility.

In the current study, we used the smoothing technique of Lapata, Keller, and McDonald (2001) to recreate not only adjective-noun bigrams, but also noun-noun and verb-object bigrams. As already mentioned in Section 1.1, it was assumed that the noun is the predicate in adjective-noun bigrams; for noun-noun bigrams, we treated the right noun as the predicate, while for verb-object bigrams, we treated the verb as the predicate. We applied Lapata, Keller, and McDonald's (2001) technique to the unseen bigrams for all three bigram types. We also used it on the seen bigrams, which were treated as unseen by removing all instances of the bigrams from the training corpus.

To test the claim that web frequencies can be used to overcome data sparseness, we correlated the frequencies recreated using class-based smoothing on the BNC with the frequencies obtained from the web. The correlation coefficients are listed in Table 12 for both seen and unseen bigrams. In all cases, a significant correlation between web counts

and recreated counts is obtained. For seen bigrams, the correlation coefficient ranged from .344 to .362 for Altavista counts, and from .330 to .349 for Google counts. For unseen bigrams, the correlations were somewhat higher, ranging from .386 to .439 for Altavista counts, and from .397 to .444 for Google counts. For both seen and unseen bigrams, there was only a very small difference between the correlation coefficients obtained with the two search engines.

It is also interesting to compare the performance of class-based smoothing and web counts on the task of predicting plausibility judgments. The correlation coefficients are listed in Table 11. The recreated frequencies are correlated significantly with all three types of bigrams, both for seen and unseen bigrams. For the seen bigrams, we found that the correlation coefficients obtained using smoothed counts were significantly lower than the upper bound for all three types of bigrams ($t(87) = 3.01, p < .01$; $t(87) = 3.23, p < .01$; $t(87) = 3.43, p < .01$). This also held for the unseen bigrams: the correlations obtained using smoothing were significantly lower than the upper bound for all three types of bigrams ($t(87) = 1.86, p < .05$; $t(87) = 1.97, p < .05$; $t(87) = 3.36, p < .01$).

Recall that the correlation coefficients obtained using the web counts were not found to be significantly different from the upper bound, which indicates that web counts are better predictors of plausibility than smoothed counts. This fact was confirmed by further significance testing: for seen bigrams, we found that the Altavista correlation coefficients were significantly higher than correlation coefficients obtained using smoothing, for all three types of bigrams ($t(87) = 3.31, p < .01$; $t(87) = 4.11, p < .01$; $t(87) = 4.32, p < .01$). This also held for Google counts ($t(87) = 3.16, p < .01$; $t(87) = 4.02, p < .01$; $t(87) = 4.03, p < .01$). For unseen bigrams, the Altavista coefficients and the coefficients obtained using smoothing were not significantly different for adjective-noun bigrams, but the difference reached significance for noun-noun and verb-object bigrams ($t(87) = 2.08, p < .05$; $t(87) = 2.53, p < .01$). For Google counts, the difference was again not significant for adjective-noun bigrams, but it reached significance for noun-noun and verb-object bigrams ($t(87) = 2.34, p < .05$; $t(87) = 2.15, p < .05$).

Finally, we conducted a small study to investigate the validity of the counts that were recreated using class-based smoothing. We correlated the recreated counts for the seen bigrams with their actual BNC and NANTC frequencies. The correlation coefficients are reported in Tables 8 and 9. We found that the correlation between recreated counts and corpus counts was significant for all three types of bigrams, for both corpora. This demonstrates that the smoothing technique we employed generates realistic corpus counts, in the sense that the recreated counts are correlated with the actual counts. However, the correlation coefficients obtained using web counts were always substantially higher than those obtained using smoothed counts. These differences were significant for the BNC counts for Altavista ($t(87) = 8.38, p < .01$; $t(87) = 5.00, p < .01$; $t(87) = 5.03, p < .01$) and Google ($t(87) = 8.35, p < .01$; $t(87) = 5.00, p < .01$; $t(87) = 5.03, p < .01$). They were also significant for the NANTC counts for Altavista ($t(87) = 4.12, p < .01$; $t(87) = 3.72, p < .01$; $t(87) = 6.58, p < .01$) and Google ($t(87) = 4.08, p < .01$; $t(87) = 3.06, p < .01$; $t(87) = 6.47, p < .01$).

To summarize, the results presented in this section indicate that web counts are indeed a valid way of obtaining counts for bigrams that are unseen in a given corpus: they correlate reliably with counts recreated using class-based smoothing. For seen bigrams, we found that web counts correlate with counts that were recreated using smoothing techniques (after removing the seen bigrams from the training corpus). For the task of predicting plausibility judgments, we were able to show that web counts outperform recreated counts, both for seen and for unseen bigrams. Finally, we found that web counts for seen bigrams correlate better than recreated counts with the real corpus counts.

It is beyond the scope of the present study to undertake a full comparison between

web counts and frequencies recreated using all available smoothing techniques (and all available taxonomies that might be used for class-based smoothing). The smoothing method discussed above is simply one type of class-based smoothing. Other more sophisticated class-based methods do away with the simplifying assumption that the argument co-occurring with a given predicate (adjective, noun, verb) is distributed evenly across its conceptual classes and attempt to find the right level of generalization in a concept hierarchy, by discounting for example the contribution of very general classes (Clark and Weir, 2001; McCarthy, 2000; Li and Abe, 1998). Other smoothing approaches such as discounting (Katz, 1987) and distance-weighted averaging (Grishman and Sterling, 1994; Dagan, Lee, and Pereira, 1999) recreate counts of unseen word combinations by exploiting only corpus-internal evidence, without relying on taxonomic information. Our goal was to demonstrate that frequencies retrieved from the web are a viable alternative to conventional smoothing methods when data is sparse; we do not claim that our web-based method is necessarily superior to smoothing or that it should be generally preferred over smoothing methods. However, the next section will present a small-scale study that compares the performance of several smoothing techniques with the performance of web counts on a standard task from the literature.

2.4 Pseudo-disambiguation

In the smoothing literature, recreated frequencies are typically evaluated using pseudo-disambiguation (Clark and Weir, 2001; Dagan, Lee, and Pereira, 1999; Lee, 1999; Pereira, Tishby, and Lee, 1993; Prescher, Riezler, and Rooth, 2000; Rooth et al., 1999). The aim of this task is to decide if a given algorithm recreates frequencies that make it possible to distinguish between seen and unseen bigrams in a given corpus. A set of pseudo-bigrams is constructed according to a set of criteria (detailed below) that ensure that they are unattested in the training corpus. Then the seen bigrams are removed from the training data and the smoothing method is used to recreate the frequencies of both the seen bigrams and the pseudo-bigrams. The smoothing method is then evaluated by comparing the frequencies it recreates for both types of bigrams.

We evaluated our web counts by applying the pseudo-disambiguation procedure that Rooth et al. (1999), Prescher, Riezler, and Rooth (2000), and Clark and Weir (2001) employed for evaluating recreated verb-object bigram counts. Take a verb-object bigram (v, n) that is seen in a given corpus and pair the noun n with a randomly chosen verb v' that fails to take n as its object. This results in an unseen verb-object bigram (v', n) . Now treat the seen bigram as unseen (i.e., remove all of its occurrences from the training corpus) and recreate the frequencies of both the seen and the unseen bigram (using smoothing, or web counts, in our case). The task is then to decide which of the two verbs v and v' take the noun n as their object. For this, the recreated bigram frequency is used: the bigram with the higher recreated frequency (or probability) is taken to be the seen bigram. If this bigram is really the seen one, then the disambiguation is correct. The overall percentage of correct disambiguations is a measure of the quality of the recreated frequencies (or probabilities). In the following, we will first describe in some detail the experiments that Rooth et al. (1999) and Clark and Weir (2001) conducted. We will then discuss how we replicated their experiments using the web as an alternative smoothing method.

Rooth et al. (1999) use pseudo-disambiguation to evaluate a class-based model that is derived from unlabeled data using the expectation maximization (EM) algorithm. From a data set of 1,280,712 (v, n) pairs (obtained from the BNC using Carroll and Rooth's (1998) parser), they randomly selected 3,000 pairs, such that each pair contained a fairly frequent verb and noun (only verbs and nouns that occurred between 30 and 3,000 times in the data were considered). For each pair (v, n) a fairly frequent verb v' was randomly

Table 13

Percentage of correct disambiguations on the pseudo-disambiguation task using web counts and counts recreated using EM-based clustering (Rooth et al., 1999)

Data set	N	Altavista cond. prob.	Altavista joint prob.	Rooth et al.
Subject	717	71.2	68.5	–
Objects	620	85.2	77.5	–
Subjects and objects	1337	77.7	72.7	80.0

chosen such that the pair (v', n) did not occur in the data set. Given the set of (v, n, v') triples (a total of 1,337), the task was to decide whether (v, n) or (v', n) is the correct (i.e., seen) pair by comparing the probabilities $P(n|v)$ and $P(n|v')$. The probabilities were recreated using Rooth et al.’s (1999) EM-based clustering model on a training set from which all seen pairs (v, n) had been removed. An accuracy of 80% on the pseudo-disambiguation task was achieved (see Table 13).

Prescher, Riezler, and Rooth (2000) evaluate Rooth et al.’s (1999) EM-based clustering model again using pseudo-disambiguation, but on a separate data set using a slightly different method for constructing the pseudo-bigrams. They used a set of 298 (v, n, n') BNC triples where (v, n) was chosen as in Rooth et al. (1999) but paired with a randomly chosen noun n' . Given the set of (v, n, n') triples the task was to decide whether (v, n) or (v, n') is the correct pair. Prescher, Riezler, and Rooth (2000) report pseudo-disambiguation results with two clustering models: (a) Rooth et al.’s (1999) clustering approach that models the semantic fit between a verb and its argument (VA model) and (b) a refined version of this approach that only models the fit between a verb and its object (VO model), disregarding other arguments of the verb. The results of the two models on the pseudo-disambiguation task are shown in Table 14.

At this point, it is important to note that both Rooth et al. (1999) and Prescher, Riezler, and Rooth (2000) do not use pseudo-disambiguation for the final evaluation of their models. Rather, the performance on the pseudo-disambiguation task is used to optimize the model parameters. The results in Tables 13 and 14 list the pseudo-disambiguation performance achieved for the best parameter settings. In other words, these results were obtained on the development set (i.e., on the same data set that was used to optimize the parameters), not on a completely unseen test set. This procedure is well-justified in the context of Rooth et al.’s (1999) and Prescher, Riezler, and Rooth’s (2000) work, which aims at building models of lexical semantics, not of pseudo-disambiguation. Therefore, they carried out their final evaluations on unseen test sets for the tasks of lexicon induction (Rooth et al., 1999) and target language disambiguation (Prescher, Riezler, and Rooth, 2000), once the model parameters had been fixed using the pseudo-disambiguation development set.⁸

Clark and Weir (2002) use a setting that that is similar to that of Rooth et al. (1999) and Prescher, Riezler, and Rooth (2000); here pseudo-disambiguation is employed

⁸ Stefan Riezler (personal communication, 2003) points out that the main variance in Rooth et al.’s (1999) pseudo-disambiguation results comes from the class cardinality parameter (start values only account for 2% of the performance and iterations do not seem to make a difference at all). Figure 3 of Rooth et al. (1999) shows that a performance of more than 75% is obtained for every reasonable choice of classes. This indicates that a “proper” pseudo-disambiguation setting with separate development and test data would have resulted in a similar choice of class cardinality, and thus achieved the same 80% performance that is cited in Table 13.

Table 14

Percentage of correct disambiguations on the pseudo-disambiguation task using web counts and counts recreated using EM-based clustering (Prescher, Riezler, and Rooth, 2000)

Data set	N	Altavista cond. prob.	Altavista joint prob.	VA model	VO model
Subjects	159	66.7	59.1	–	–
Objects	139	70.5	66.2	–	–
Subjects and objects	298	68.5	62.4	79.0	88.3

Table 15

Percentage of correct disambiguations on the pseudo-disambiguation task using web counts and counts recreated using class-based smoothing (Clark and Weir, 2002)

Data set	N	Altavista cond. prob.	Altavista joint prob.	Clark & Weir	Li & Abe	Resnik
Objects (low freq.)	3000	83.9	81.1	72.4	62.9	62.6
Objects (high freq.)	3000	87.7	85.3	73.9	68.3	63.9

to evaluate the performance of a class-based probability estimation method. In order to address the problem of estimating conditional probabilities in the face of sparse data, Clark and Weir (2002) define probabilities in terms of classes in a semantic hierarchy and propose hypothesis testing as a means of determining a suitable level of generalization in the hierarchy. Clark and Weir (2002) report pseudo-disambiguation results on two data sets, with an experimental setup similar to Rooth et al. (1999). For the first data set, 3,000 pairs were randomly chosen from 1.3 million (v, n) tuples extracted from the BNC (using the parser of Briscoe and Carroll (1997)). The selected pairs contained relatively frequent verbs (occurring between 500 and 5,000 times in the data). The data sets were constructed as proposed by Rooth et al. (1999). The procedure for creating the second data set was identical, but this time only verbs that occurred between 100 and 1,000 times were considered. Clark and Weir (2002) further compared their approach against Resnik's (1993) selectional association model and Li and Abe's (1998) tree cut models on the same data sets. These methods are directly comparable as they can be used for class-based probability estimation and address the question of how to find a suitable level of generalization in a hierarchy (i.e., WordNet). The results of the three methods on the two data sets are shown in Table 15.

We employed the same pseudo-disambiguation method to test whether web-based frequencies can be used for distinguishing between seen and artificially constructed unseen bigrams. We obtained the data sets of Rooth et al. (1999), Prescher, Riezler, and Rooth (2000), and Clark and Weir (2002) described above. Given a set of (v, n, v') triples, the task was to decide whether (v, n) or (v', n) was the correct pair. We obtained Altavista counts for $f(v, n)$, $f(v', n)$, and $f(n)$ as described in Section 1.3.⁹ Then we used two models for pseudo-disambiguation: the joint probability model compared the joint probability estimates $f(v, n)$ and $f(v', n)$ and predicted that the bigram with the highest

⁹ We only used Altavista counts, as there was virtually no difference between Altavista and Google counts in our previous evaluations (see Sections 2.1–2.3). Google only allows 1000 queries per day (for registered users), which makes it time-consuming to obtain large numbers of Google counts. Altavista has no such restriction.

estimate is the seen one. The conditional probability model compared the conditional probability estimates $f(v, n)/f(v)$ and $f(v', n)/f(v')$ and again selected as the seen bigram the one with the highest estimate (in both cases, ties were resolved by choosing at random).¹⁰ The same two models were used to perform pseudo-disambiguation for the (v, n, n') triples, where we have to choose between (v, n) and (v, n') . Here, the probability estimates $f(v, n)$ and $f(v, n')$ were used for the joint probability model, and $f(v, n)/f(n)$ and $f(v, n')/f(n')$ for the conditional probability model.

The results for Rooth et al.'s (1999) data set are given in Table 13. The conditional probability model achieves a performance of 71.2% correct for subjects and 85.2% correct for objects. The performance on the whole data set is 77.7%, which is below the performance of 80.0% reported by Rooth et al. (1999). However, the difference is not significant using a χ^2 -test comparing the number of correct and incorrect classifications ($\chi^2(1) = 2.02, p = .16$). The joint probability model performs consistently worse than the conditional probability model, it achieves an overall accuracy of 72.7%, which is significantly lower than the accuracy of the Rooth et al. (1999) model ($\chi^2(1) = 19.50, p < .01$).

A similar picture emerges on Prescher, Riezler, and Rooth's (2000) data set (see Table 14). The conditional probability model achieves an accuracy of 66.7% for subjects and 70.5% for objects. The combined performance of 68.5% is significantly lower than the performance of both the VA model ($\chi^2(1) = 7.78, p < .01$) and the VO model ($\chi^2(1) = 33.28, p < .01$) reported by Prescher, Riezler, and Rooth (2000). Again, the joint probability model performs worse than the conditional probability model, achieving an overall accuracy of 62.4%.

We also applied our web-based method to the pseudo-disambiguation data set of Clark and Weir (2002). Here, the conditional probability model reached a performance of 83.9% correct on the low frequency data set. This is significantly higher than the highest performance of 72.4% reported by Clark and Weir (2002) on the same data set ($\chi^2(1) = 115.50, p < .01$). The joint probability model performs worse than the conditional model, at 81.1%. However, this is still significantly better than the best result of Clark and Weir (2002) ($\chi^2(1) = 63.14, p < .01$). The same pattern is observed for the high frequency data set, where the conditional probability model achieves 87.7% correct, and thus significantly outperforms Clark and Weir (2002), who obtained 73.9% ($\chi^2(1) = 283.73, p < .01$). The joint probability model achieved 85.3% on this data set, also significantly outperforming Clark and Weir (2002) ($\chi^2(1) = 119.35, p < .01$).

To summarize, we demonstrated that the simple web-based approach proposed in this paper yields results for pseudo-disambiguation that outperform class-based smoothing techniques, such as the ones proposed by Resnik (1993), Li and Abe (1998), and Clark and Weir (2002). We were also able to show that a web-based approach is able to achieve the same performance as an EM-based smoothing model proposed by Rooth et al.'s (1999). However, the web-based approach was not able to outperform the more sophisticated EM-based model of Prescher, Riezler, and Rooth (2000). Another result we obtained was that web-based models that uses joint probabilities (where unigram frequencies are used to normalize the bigram frequencies) generally outperforms a more simple-minded approach that relies directly on bigram frequencies for pseudo-disambiguation.

There are a number of reasons why our results regarding pseudo-disambiguation have to be treated with some caution. First of all, the two smoothing methods (i.e., EM-based clustering and class-based probability estimation using WordNet) were not evaluated on

¹⁰ The probability estimates are $P(a, b) = f(a, b)/N$ and $P(a|b) = f(a, b)/f(b)$ for the joint probability and the conditional probability, respectively. However, the corpus size N can be ignored, as it is constant.

the same data set, and therefore the two results are not directly comparable. For instance, Clark and Weir's (2002) data set is substantially less noisy than Rooth et al.'s (1999) and Prescher, Riezler, and Rooth's (2000), as it only contains words and nouns that occur in WordNet. Furthermore, Stephen Clark (personal communication, 2003) points out that WordNet-based approaches are at a disadvantage when it comes to pseudo-disambiguation. Pseudo-disambiguation assumes that the correct pair is unseen in the training data; this makes the task deliberately hard because some of the pairs might be frequent enough to obtain reliable corpus counts without having to use WordNet (using WordNet is likely to be more noisy than using the actual counts). Another problem with WordNet-based approaches is that they offer no systematic treatment of word sense ambiguity, which puts them at a disadvantage with respect to approaches that do not rely on a predefined inventory of word senses.

Finally, recall that the results for the EM-based approaches in Tables 13 and 14 were obtained on the development set (as pseudo-disambiguation was used as a means of parameter tuning by Rooth et al. (1999) and Prescher, Riezler, and Rooth (2000)). It is possible that this fact inflates the performance values for the EM-based approaches (but see Footnote 8 above).

3 Conclusions

This paper explored a novel approach to overcoming data sparseness. If a bigram is unseen in a given corpus, conventional approaches recreate its frequency using techniques such as back-off, linear interpolation, class-based smoothing or distance-weighted averaging (see Dagan, Lee, and Pereira (1999) and Lee (1999) for overviews). The approach proposed here does not recreate the missing counts, but instead retrieves them from a corpus that is much larger (but also much more noisy) than any existing corpus: it launches queries to a search engine in order to determine how often a bigram occurs on the web.

We systematically investigated the validity of this approach by using it to obtain frequencies for predicate-argument bigrams (adjective-noun, noun-noun, and verb-object bigrams). We first applied the approach to seen bigrams randomly sampled from the BNC. We found that the counts obtained from the web are highly correlated with the counts obtained from the BNC. We then obtained bigram counts from NANTC, a corpus that is substantially larger than the BNC. Again, we found that web counts are highly correlated with corpus counts. This indicates that web queries can generate frequencies that are comparable to the ones obtained from a balanced, carefully edited corpus such as the BNC, but also from a large news text corpus such as NANTC.

Secondly, we performed an evaluation that used the web frequencies to predict human plausibility judgments for predicate-argument bigrams. The results show that web counts correlate reliably with judgments, for all three types of predicate-argument bigrams tested, both seen and unseen. For the seen bigrams, we showed that the web frequencies correlate better with judged plausibility than corpus frequencies.

To substantiate the claim that the web counts can be used to overcome data sparseness, we compared bigram counts obtained from the web with bigram counts recreated using a class-based smoothing technique (a variant of the one proposed by Resnik (1993)). We found that web frequencies and recreated frequencies are reliably correlated, and that web frequencies are better at predicting plausibility judgments than smoothed frequencies. This holds both for unseen bigrams and for seen bigrams that are treated as unseen by omitting them from the training corpus.

Finally, we tested the performance of our frequencies in a standard pseudo-disambiguation task. We applied our method to three data sets from the literature. The results show that web counts outperform counts recreated using a number of class-

based smoothing techniques. However, counts recreated using an EM-based smoothing approach yielded better pseudo-disambiguation performance than web counts.

To summarize, we have proposed a simple heuristic method for obtaining bigram counts from the web. Using four different types of evaluation, we demonstrated that this simple heuristic method is sufficient to obtain valid frequency estimates. It seems that the large amount of data available outweighs the problems associated with using the web as a corpus (such as the fact that it is noisy and unbalanced).

A number of questions arise for future research: (a) Are web frequencies suitable for probabilistic modeling, in particular since web counts are not perfectly normalized as Zhu and Rosenfeld (2001) have shown? (b) How can existing smoothing methods benefit from web counts? (c) How do the results reported in this paper carry over to languages other than English (for which much less web data is available)? (d) What is the effect of the noise introduced by our heuristic approach? The last question could be assessed by reproducing our results using a snapshot of the web, on which argument relations can be extracted more accurately using POS tagging and chunking techniques.

Finally, it will be crucial to test the usefulness of web-based frequencies for realistic NLP tasks. Preliminary results are reported by Lapata and Keller (2003), who use web counts successfully for a range of NLP tasks, including candidate selection for machine translation, context-sensitive spelling correction, bracketing and interpretation of compounds, adjective ordering, and PP-attachment.

Acknowledgments

This work was conducted while both authors were at the Department of Computational Linguistics, Saarland University, Saarbrücken. The work was inspired by a talk that Gregory Grefenstette gave in Saarbrücken in 2001 about his research on using the web as a corpus. The present paper is an extended and revised version of Keller, Lapata, and Ourioupina (2002). Stephen Clark and Stefan Riezler provided valuable comments on this research. We are also grateful to four anonymous reviewers for *Computational Linguistics*; their feedback helped to substantially improve the present paper. Special thanks are due to Stephen Clark and Detlef Prescher for making their pseudo-disambiguation data sets available.

References

- Abney, Steve. 1996. Partial parsing via finite-state cascades. In John Carroll, editor, *Workshop on Robust Parsing*, pages 8–15, 8th European Summer School in Logic, Language and Information, Prague, Czech Republic.
- Abney, Steve and Marc Light. 1999. Hiding a semantic class hierarchy in a Markov model. In Andrew Kehler and Andreas Stolcke, editors, *Proceedings of the ACL Workshop on Unsupervised Learning in Natural Language Processing*, pages 1–8, College Park, Maryland.
- Agirre, Eneko and David Martinez. 2000. Exploring automatic word sense disambiguation with decision lists and the web. In *Proceedings of the 18th International Conference on Computational Linguistics*, pages 11–19, Saarbrücken, Germany.
- Banko, Michele and Eric Brill. 2001a. Mitigating the paucity-of-data problem: Exploring the effect of training corpus size on classifier performance for natural language processing. In James Allan, editor, *Proceedings of the 1st International Conference on Human Language Technology Research*, San Francisco, California. Morgan Kaufmann.
- Banko, Michele and Eric Brill. 2001b. Scaling to very very large corpora for natural language disambiguation. In *Proceedings of the 39th Annual Meeting of the Association for Computational Linguistics and the 10th Conference of the European Chapter of the Association for Computational Linguistics*, pages 26–33, Toulouse, France.
- Bard, Ellen Gurman, Dan Robertson, and Antonella Sorace. 1996. Magnitude estimation of linguistic acceptability. *Language*, 72(1):32–68.
- Briscoe, Ted and John Carroll. 1997. Automatic extraction of subcategorization from corpora. In *Proceedings of the 5th Conference on Applied Natural Language Processing*, pages 356–363, Washington, DC.

- Burnard, Lou, 1995. *Users Reference Guide, British National Corpus*. British National Corpus Consortium, Oxford University Computing Services, Oxford, England.
- Carroll, Glenn and Mats Rooth. 1998. Valence induction with a head-lexicalized PCFG. In Nancy Ide and Atro Voutilainen, editors, *Proceedings of the 3rd Conference on Empirical Methods in Natural Language Processing*, pages 36–45, Granada, Spain.
- Clark, Stephen and David Weir. 2001. Class-based probability estimation using a semantic hierarchy. In *Proceedings of the 2nd Conference of the North American Chapter of the Association for Computational Linguistics*, pages 95–102, Pittsburgh, Pennsylvania.
- Clark, Stephen and David Weir. 2002. Class-based probability estimation using a semantic hierarchy. *Computational Linguistics*, 28(2):187–206.
- Corley, Martin and Christoph Scheepers. 2002. Syntactic priming in English sentence production: Categorical and latency evidence from an internet-based study. *Psychonomic Bulletin and Review*, 9(1):126–131.
- Corley, Steffan, Martin Corley, Frank Keller, Matthew W. Crocker, and Shari Trewin. 2001. Finding syntactic structure in unparsed corpora: The Gsearch corpus query system. *Computers and the Humanities*, 35(2):81–94.
- Cowart, Wayne. 1997. *Experimental Syntax: Applying Objective Methods to Sentence Judgments*. Sage Publications, Thousand Oaks, California.
- Curran, James. 2002. Scaling context space. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 231–238, Philadelphia, Pennsylvania.
- Dagan, Ido, Lillian Lee, and Fernando Pereira. 1999. Similarity-based models of word cooccurrence probabilities. *Machine Learning*, 34(1):43–69.
- Grefenstette, Gregory. 1994. *Explorations in Automatic Thesaurus Discovery*. Kluwer Academic Publishers, Boston, Massachusetts.
- Grefenstette, Gregory. 1998. The World Wide Web as a resource for example-based machine translation tasks. In *Proceedings of the ASLIB Conference on Translating and the Computer*, London, England.
- Grefenstette, Gregory and Jean Nioche. 2000. Estimation of English and non-English language use on the WWW. In *Proceedings of the RIAO Conference on Content-Based Multimedia Information Access*, pages 237–246, Paris, France.
- Grishman, Ralph and John Sterling. 1994. Generalizing automatically generated selectional patterns. In *Proceedings of the 15th International Conference on Computational Linguistics*, pages 742–747, Kyoto, Japan.
- Hindle, Donald and Mats Rooth. 1993. Structural ambiguity and lexical relations. *Computational Linguistics*, 19(1):103–120.
- Jones, Rosie and Rayid Ghani. 2000. Automatically building a corpus for a minority language from the web. In *Proceedings of the Student Research Workshop at the 38th Annual Meeting of the Association for Computational Linguistics*, pages 29–36, Hong Kong.
- Karp, Daniel, Yves Schabes, Martin Zaidel, and Dania Egedi. 1992. A freely available wide coverage morphological analyzer for English. In *Proceedings of the 14th International Conference on Computational Linguistics*, pages 950–954, Nantes, France.
- Katz, Slava M. 1987. Estimation of probabilities from sparse data for the language model component of a speech recognizer. *IEEE Transactions on Acoustics Speech and Signal Processing*, 33(3):400–401.
- Keller, Frank and Theodora Alexopoulou. 2001. Phonology competes with syntax: Experimental evidence for the interaction of word order and accent placement in the realization of information structure. *Cognition*, 79(3):301–372.
- Keller, Frank and Ash Asudeh. 2001. Constraints on linguistic coreference: Structural vs. pragmatic factors. In Johanna D. Moore and Keith Stenning, editors, *Proceedings of the 23rd Annual Conference of the Cognitive Science Society*, pages 483–488, Mahwah, New Jersey. Lawrence Erlbaum Associates.
- Keller, Frank, Martin Corley, Steffan Corley, Lars Konieczny, and Amalia Todirascu. 1998. WebExp: A Java toolbox for web-based psychological experiments. Technical Report HCRC/TR-99, Human Communication Research Centre, University of Edinburgh, Edinburgh, Scotland.
- Keller, Frank, Maria Lapata, and Olga Ourioupina. 2002. Using the web to overcome data sparseness. In Jan Hajič and Yuji Matsumoto, editors, *Proceedings*

- of the *Conference on Empirical Methods in Natural Language Processing*, pages 230–237, Philadelphia, Pennsylvania.
- Lapata, Maria. 2001. A corpus-based account of regular polysemy: The case of context-sensitive adjectives. In *Proceedings of the 2nd Conference of the North American Chapter of the Association for Computational Linguistics*, pages 63–70, Pittsburgh, Pennsylvania.
- Lapata, Maria. 2002. The disambiguation of nominalisations. *Computational Linguistics*, 28(3):357–388.
- Lapata, Maria and Frank Keller. 2003. Evaluating the performance of unsupervised web-based models for a range of NLP tasks. Unpublished manuscript, University of Sheffield, Sheffield, England and University of Edinburgh, Edinburgh, Scotland.
- Lapata, Maria, Frank Keller, and Scott McDonald. 2001. Evaluating smoothing algorithms against plausibility judgments. In *Proceedings of the 39th Annual Meeting of the Association for Computational Linguistics and the 10th Conference of the European Chapter of the Association for Computational Linguistics*, pages 346–353, Toulouse, France.
- Lapata, Maria, Scott McDonald, and Frank Keller. 1999. Determinants of adjective-noun plausibility. In *Proceedings of the 9th Conference of the European Chapter of the Association for Computational Linguistics*, pages 30–36, Bergen, Norway.
- Lauer, Mark. 1995. *Designing Statistical Language Learners: Experiments on Compound Nouns*. Ph.D. thesis, Macquarie University, Sydney, Australia.
- Lee, Lillian. 1999. Measures of distributional similarity. In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics*, pages 25–32, College Park, Maryland.
- Leech, Geoffrey, Roger Garside, and Michael Bryant. 1994. The tagging of the British national corpus. In *Proceedings of the 15th International Conference on Computational Linguistics*, pages 622–628, Kyoto, Japan.
- Levi, Judith N. 1978. *The Syntax and Semantics of Complex Nominals*. Academic Press, New York, New York.
- Li, Hang and Naoki Abe. 1998. Generalizing case frames using a thesaurus and the MDL principle. *Computational Linguistics*, 24(2):217–244.
- Lin, Dekang. 1994. Principar – an efficient, broad-coverage, principle-based parser. In *Proceedings of the 15th International Conference on Computational Linguistics*, pages 482–488, Kyoto, Japan.
- Lin, Dekang. 1998. Dependency-based evaluation of MINIPAR. In *Proceedings on of the LREC Workshop on the Evaluation of Parsing Systems*, pages 48–56, Granada, Spain.
- Lin, Dekang. 2001. LaTaT: Language and text analysis tools. In *Proceedings of the 1st International Conference on Human Language Technology Research*, San Francisco, California. Morgan Kaufmann.
- Markert, Katja, Malvina Nissim, and Natalia N. Modjeska. 2003. Using the web for nominal anaphora resolution. In *Proceedings of the EACL Workshop on the Computational Treatment of Anaphora*, Budapest, Hungary.
- McCarthy, Diana. 2000. Using semantic preferences to identify verbal participation in role switching alternations. In *Proceedings of the 1st Conference of the North American Chapter of the Association for Computational Linguistics*, pages 256–263, Seattle, Washington.
- Mihalcea, Rada and Dan Moldovan. 1999. A method for word sense disambiguation of unrestricted text. In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics*, pages 152–158, College Park, Maryland.
- Miller, George A., 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.
- Pereira, Fernando, Naftali Tishby, and Lillian Lee. 1993. Distributional clustering of English words. In *Proceedings of the 31st Annual Meeting of the Association for Computational Linguistics*, pages 183–190, Columbus, Ohio.
- Pollard, Carl and Ivan A. Sag. 1994. *Head-Driven Phrase Structure Grammar*. University of Chicago Press, Chicago, Illinois.
- Prescher, Detlef, Stefan Riezler, and Mats Rooth. 2000. Using a probabilistic class-based lexicon for lexical ambiguity resolution. In *Proceedings of the 18th International Conference on Computational Linguistics*, pages 649–655, Saarbrücken, Germany.

- Ratnaparkhi, Adwait. 1998. Unsupervised statistical models for prepositional phrase attachment. In *Proceedings of the 17th International Conference on Computational Linguistics and 36th Annual Meeting of the Association for Computational Linguistics*, pages 1079–1085, Montréal, Canada.
- Resnik, Philip. 1999. Mining the web for bilingual text. In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics*, pages 527–534, College Park, Maryland.
- Resnik, Philip. 2000. Measuring verb similarity. In Lila R. Gleitman and Aravid K. Joshi, editors, *Proceedings of the 22nd Annual Conference of the Cognitive Science Society*, pages 399–404, Mahwah, New Jersey. Lawrence Erlbaum Associates.
- Resnik, Philip Stuart. 1993. *Selection and Information: A Class-Based Approach to Lexical Relationships*. Ph.D. thesis, University of Pennsylvania, Philadelphia, Pennsylvania.
- Rooth, Mats, Stefan Riezler, Detlef Prescher, Glenn Carroll, and Franz Beil. 1999. Inducing a semantically annotated lexicon via EM-based clustering. In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics*, pages 104–111, College Park, Maryland.
- Sampson, Geoffrey. 1995. *English for the Computer: The SUSANNE Corpus and Analytic Scheme*. Clarendon Press, Oxford, England.
- Schütze, Carson T. 1996. *The Empirical Base of Linguistics: Grammaticality Judgments and Linguistic Methodology*. University of Chicago Press, Chicago, Illinois.
- Stevens, S. S. 1975. *Psychophysics: Introduction to its Perceptual, Neural, and Social Prospects*. John Wiley, New York, New York.
- Volk, Martin. 2001. Exploiting the WWW as a corpus to resolve PP attachment ambiguities. In Paul Rayson, Andrew Wilson, Tony McEnery, Andrew Hardie, and Shereen Khoja, editors, *Proceedings of the Corpus Linguistics Conference*, pages 601–606, Lancaster, England.
- Weiss, Sholom M. and Casimir A. Kulikowski. 1991. *Computer Systems that Learn: Classification and Prediction Methods from Statistics, Neural Nets, Machine Learning, and Expert Systems*. Morgan Kaufmann, San Mateo, California.
- Zhu, Xiaojin and Ronald Rosenfeld. 2001. Improving trigram language modeling with the world wide web. In *In Proceedings of International Conference on Acoustics Speech and Signal Processing*, Salt Lake City, Utah.