

Using Semantically Motivated Estimates to Help Subcategorization Acquisition

Anna Korhonen

Computer Laboratory, University of Cambridge
Pembroke Street, Cambridge CB2 3QG, UK
alk23@cl.cam.ac.uk

Abstract

Research into the automatic acquisition of subcategorization frames from corpora is starting to produce large-scale computational lexicons which include valuable frequency information. However, the accuracy of the resulting lexicons shows room for improvement. One source of error lies in the lack of accurate back-off estimates for subcategorization frames, delimiting the performance of statistical techniques frequently employed in verbal acquisition. In this paper, we propose a method of obtaining more accurate, semantically motivated back-off estimates, demonstrate how these estimates can be used to improve the learning of subcategorization frames, and discuss using the method to benefit large-scale lexical acquisition.

1 Introduction

Manual development of large subcategorised lexicons has proved difficult because predicates change behaviour between sublanguages, domains and over time. Yet parsers depend crucially on such information, and probabilistic parsers would greatly benefit from accurate information concerning the relative frequency of different subcategorization frames (SCFs) for a given predicate.

Over the past years acquiring subcategorization dictionaries from textual corpora has become increasingly popular (e.g. Brent, 1991, 1993; Ushioda *et al.*, 1993; Briscoe and Carroll, 1997; Manning, 1993; Carroll and Rooth 1998; Gahl, 1998; Lapata, 1999, Sarkar and Zeman, 2000). The different approaches vary according to the methods used and the number of SCFs being extracted. Regardless of this, there is a ceiling on the performance of these systems at around 80% token recall¹.

¹Token recall is the percentage of SCF tokens in a sample of manually analysed text that were correctly acquired by the system.

One significant source of error lies in the statistical filtering methods frequently used to remove noise from automatically acquired SCFs. These methods are reported to be particularly unreliable for low frequency SCFs (Brent, 1991, 1993; Briscoe and Carroll, 1997; Manning, 1993; Manning and Schütze, 1999; Korhonen, Gorrell and McCarthy, 2000), resulting in a poor overall performance.

According to Korhonen, Gorrell and McCarthy (2000), the poor performance of statistical filtering can be largely explained by the zipfian nature of the data, coupled with the fact that many statistical tests are based on the assumption of two zipfian distributions correlating: the conditional SCF distribution of an individual verb ($p(scf_i|verb_j)$) and the unconditional SCF distribution of all verbs in general ($p(scf_i)$). Contrary to this assumption, however, there is no significant correlation between the two distributions.

Korhonen, Gorrell and McCarthy (2000) have showed that a simple method of filtering SCFs on the basis of their relative frequency performs more accurately than statistical filtering. This method sensitive to the sparse data problem is best integrated with smoothing. Yet the performance of the sophisticated smoothing techniques which back-off to an unconditional distribution also suffer from the lack of correlation between $p(scf_i|verb_j)$ and $p(scf_i)$.

In this paper, we propose a method for obtaining more accurate back-off estimates for SCF acquisition. Taking Levin's verb classification (Levin, 1993) as a starting point, we show that in terms of SCF distributions, individual verbs correlate better with other semantically similar verbs than with all verbs in general. On the basis of this observation, we propose classifying verbs according to their semantic class and using the conditional SCF distributions of a few other members in the

same class as back-off estimates of the class ($p(scf_i|semantic\ class_j)$).

Adopting the SCF acquisition system of Briscoe and Carroll (1997) we report an experiment which demonstrates how these estimates can be used in filtering. This is done by acquiring the conditional SCF distributions for selected test verbs, smoothing these distributions with the unconditional distribution of the respective verb class, and applying a simple method for filtering the resulting set of SCFs. Our results show that the proposed method improves the acquisition of SCFs significantly. We discuss how this method can be used to benefit large-scale SCF acquisition.

We begin by reporting our findings that the SCF distributions of semantically similar verbs correlate well (section 2). We then introduce the method we adopted for constructing the back-off estimates for the data used in our experiment (section 3.1), summarise the main features of the SCF acquisition approach (section 3.2), and describe the smoothing techniques adopted (section 3.3). Finally, we review the empirical evaluation (section 4) and discuss directions for future work (section 5).

2 Examining SCF Correlation between Semantically Similar Verbs

To examine the degree of SCF correlation between semantically similar verbs, we took Levin’s verb classification (1993) as a starting point. Levin verb classes are based on the ability of a verb to occur in specific diathesis alternations, i.e. specific pairs of syntactic frames which are assumed to be meaning preserving. The classification provides semantically-motivated sets of syntactic frames associated with individual classes.

While Levin’s shows that there is correlation between the SCFs related to the verb *sense*, our aim is to examine whether there is also correlation between the SCFs specific to the verb *form*. Unlike Levin, we are concerned with polysemic SCF distributions involving all senses of verbs. In addition, we are not only interested in the degree of correlation between sets of SCFs, but also in comparing the ranking of SCFs between distributions. Nevertheless, Levin classes provide us a useful starting point.

Focusing on five broad Levin classes - change of possession, assessment, killing, motion, and destroy verbs - we chose four test

verbs from each class and examined the degree with which the SCF distribution for these verbs correlates with the SCF distributions for two other verbs from the same Levin class. The latter verbs were chosen so that one of the verbs is a synonym, and the other a hypernym, of a test verb. We used WordNet (Miller *et al.*, 1990) for defining and recognising these semantic relations. We defined a hypernym as a test verb’s hypernym in WordNet, and a synonym as a verb which, in WordNet, shares this same hypernym with a test verb. We also examined how well the SCF distribution for the different test verbs correlates with the SCF distribution of all English verbs in general and with that of a semantically different verb (i.e. a verb belonging to a different Levin class).

We used two methods for obtaining the SCF distributions. The first was to acquire an unfiltered subcategorization lexicon for 20 million words of the British National Corpus (BNC) (Leech, 1992) data using Briscoe and Carroll’s (1997) system (introduced in section 3.2). This gives us the *observed* distribution of SCFs for individual verbs and that for all verbs in the BNC data. The second method was to manually analyse around 300 occurrences of each test verb in the BNC data. This gives us an estimate of the *correct* SCF distributions for the individual verbs. The estimate for the correct distribution of SCFs over all English verbs was obtained by extracting the number of verbs which are members of each SCF class in the ANLT dictionary (Boguraev *et al.*, 1987).

The degree of correlation was examined by calculating the Kullback-Leibler distance (KL) (Cover and Thomas, 1991) and the Spearman rank correlation coefficient (RC) (Spearman, 1904) between the different distributions².

The results given in tables 1 and 2 were obtained by correlating the observed SCF distributions from the BNC data. Table 1 shows an example of correlating the SCF distribution of the motion verb *fly* against that of (i) its hypernym *move*, (ii) synonym *sail*, (iii) all verbs in general, and (iv) *agree*, which is not related semantically. The results show that the SCF distribution for *fly* clearly correlates better with the SCF distribution for *move* and *sail* than that for all verbs and *agree*. The av-

²Note that $KL \geq 0$, with KL near to 0 denoting strong association, and $-1 \leq RC \leq 1$, with RC near to 0 denoting a low degree of association and RC near to -1 and 1 denoting strong association.

		KL	RC
<i>fly</i>	<i>move</i>	0.25	0.83
<i>fly</i>	<i>sail</i>	0.62	0.61
<i>fly</i>	all verbs	2.13	0.51
<i>fly</i>	<i>agree</i>	2.27	0.12

Table 1: Correlating the SCF distribution of *fly* against other SCF distributions

	KL	RC
hypernym	0.65	0.71
synonym	0.71	0.66
all verbs	1.59	0.41
semantically different verb	1.74	0.38

Table 2: Overall correlation results

erage results for all test verbs given in table 2 indicate that the degree of SCF correlation is the best with semantically similar verbs. Hypernym and synonym relations are nearly as good, the majority of verbs showing slightly better SCF correlation with hypernyms. The SCF correlation between individual verbs, and verbs in general, is poor, but still better than with semantically unrelated verbs.

These findings with the observed SCF distributions hold as well with the correct SCF distributions, as seen in table 3. The results show that in terms of SCF distributions, verbs in all classes examined correlate better with their hypernym verbs than with all verbs in general.

As one might expect, the polysemy of the individual verbs affects the degree of SCF correlation between semantically similar verbs. The degree of SCF correlation is higher with those verbs whose predominant³ sense is involved with the Levin class examined. For example, the SCF distribution for the killing verb *murder* correlates better with that for the verb *kill* than that for the verb *execute*, whose predominant sense is not involved with killing verbs.

These results show that the verb sense specific SCF correlation observed by Levin extends to the verb form specific SCF correlation and applies to the ranking of SCFs as well. This suggests that we can obtain more accurate back-off estimates for verbal acquisition by basing them on a semantic verb type. To find out whether such semantically motivated

³Predominant sense refers here to the most frequent sense of verbs in WordNet.

estimates can be used to improve SCF acquisition, we performed an experiment which we describe below.

3 Experiment

3.1 Back-off Estimates for the Data

The test data consisted of a total of 60 verbs from 12 broad Levin classes, listed in table 4. Two of the examined Levin classes were collapsed together with another similar Levin class, making the total number of test classes 10. The verbs were chosen at random, subject to the constraint that they occurred frequently enough in corpus data⁴ and when applicable, represented different sub-classes of each examined Levin class. To reduce the problem of polysemy, we required that the predominant sense of each verb corresponds to the Levin class in question. This was ensured by manually verifying that the most frequent sense of a verb in WordNet corresponds to the sense involved in the particular Levin class.

To obtain the back-off estimates, we chose 4-5 representative verbs from each verb class and obtained correct SCF distributions for these verbs by manually analysing around 300 occurrences of each verb in the BNC data. We merged the resulting set of SCF distributions to construct the unconditional SCF distribution for the verb class. This approach was taken to minimise the sparse data problem and cover SCF variations within verb classes and due to polysemy. The back-off estimates

⁴We required at least 300 occurrences for each verb. This was merely to guarantee accurate enough testing, as we evaluated our results against manual analysis of corpus data (see section 4).

Verb class	Hypernym		All Verbs	
	KL	RC	KL	RC
change of possession	0.61	0.64	1.16	0.38
assessment	0.28	0.71	0.73	0.48
killing	0.70	0.63	1.14	0.37
destroy	0.30	0.60	1.19	0.29
motion	0.29	0.73	1.72	0.42
AVERAGE	0.44	0.66	1.19	0.39

Table 3: Correlation results for five verb classes

Verb class	Test verbs
putting	<i>place, lay, drop, pour, load, fill</i>
sending and carrying, exerting force	<i>send, ship, carry, bring, transport pull, push</i>
change of possession	<i>give, lend, contribute, donate, offer provide, supply, acquire, buy</i>
assessment, searching	<i>analyse fish, explore, investigate</i>
social interaction	<i>agree, communicate, struggle, marry, meet, visit</i>
killing	<i>kill, murder, slaughter, strangle</i>
destroy	<i>demolish, destroy, ruin, devastate</i>
appearance, disappearance and occurrence	<i>arise, emerge, disappear, vanish</i>
motion	<i>arrive, depart, march, move, slide, swing travel, walk, fly, sail, dance</i>
aspectual	<i>begin, end, start, terminate, complete</i>

Table 4: Test data

for motion verbs, for example, were obtained by merging the SCF distributions of the verbs *march, move, fly, slide* and *sail*. Each verb used in obtaining the estimates was excluded when testing the verb itself. For example, when acquiring subcategorization for the verb *fly*, estimates were obtained only using verbs *march, move, slide* and *sail*.

3.2 Framework for SCF Acquisition

Briscoe and Carroll’s (1997) verbal acquisition system distinguishes 163 SCFs and returns relative frequencies for each SCF found for a given predicate. The SCFs are a superset of classes found in the ANLT and COMLEX (Grishman *et al.*, 1994) dictionaries. They incorporate information about control of predicative arguments, as well as alternations such as extraposition and particle movement. The system employs a shallow parser to obtain the subcategorization information. Potential SCF entries are filtered before the final SCF lexicon is produced. While Briscoe and Carroll (1997) used a statistical filter based on bi-

nomial hypothesis test, we adopted another method, where the conditional SCF distribution from the system is smoothed before filtering the SCFs, using the different techniques introduced in section 3.3. After smoothing, filtering is performed by applying a threshold to the resulting set of probability estimates. We used training data to find an optimal average threshold for each verb class examined. This filtering method allows us to examine the benefits of smoothing without introducing problems based on the statistical filter.

3.3 Smoothing

3.3.1 Add One Smoothing

Add one smoothing⁵ has the effect of giving some of the probability space to the SCFs unseen in the conditional distribution. As it assumes a uniform prior on events, it provides a baseline smoothing method against which the

⁵See (Manning and Schütze, 1999) for detailed information about the smoothing techniques discussed here.

more sophisticated methods can be compared. Let $c(x_n)$ be the frequency of a SCF given a verb, N the total number of SCF tokens for this verb in the conditional distribution, and C the total number of SCF types. The estimated probability of the SCF is:

$$P(x_n) = \frac{c(x_n) + 1}{N + C} \quad (1)$$

3.3.2 Katz Backing-off

In Katz backing-off (Katz, 1987), some of the probability space is given to the SCFs unseen or of low frequency in the conditional distribution. This is done by backing-off to an unconditional distribution. Let $p(x_n)$ be a probability of a SCF in the conditional distribution, and $p(x_{np})$ its probability in the unconditional distribution, obtained by maximum likelihood estimation. The estimated probability of the SCF is calculated as follows:

$$P(x_n) = \begin{cases} (1 - d) \times p(x_n) & \text{if } c(x_n) > c_1 \\ \alpha \times p(x_{np}) & \text{otherwise} \end{cases} \quad (2)$$

The cut off frequency c_1 is an empirically defined threshold determining whether to back-off or not. When counts are lower than c_1 they are held too low to give an accurate estimate, and we back-off to an unconditional distribution. In this case, we discount $p(x_n)$ a certain amount to reserve some of the probability space for unseen and very low frequency SCFs. The discount (d) is defined empirically, and α is a normalization constant which ensures that the probabilities of the resulting distribution sum to 1.

3.3.3 Linear Interpolation

While Katz backing-off consults different estimates depending on their specificity, linear interpolation makes a linear combination of them. Linear interpolation is used here for the simple task of combining a conditional distribution with an unconditional one. The estimated probability of the SCF is given by

$$P(x_n) = \lambda_1(p(x_n)) + \lambda_2(p(x_{np})) \quad (3)$$

where the λ_i denotes weights for different context sizes (obtained by optimising the smoothing performance on the training data for all x_n) and sum to 1.

4 Evaluation

4.1 Method

To evaluate the approach, we took a sample of 20 million words of the BNC and extracted all sentences containing an occurrence of one of the 60 test verbs on average of 3000 citations of each. The sentences containing these verbs were processed by the SCF acquisition system, and the smoothing methods were applied before filtering. We also obtained results for a baseline without any smoothing.

The results were evaluated against a manual analysis of the corpus data. This was obtained by analysing up to a maximum of 300 occurrences for each test verb in BNC or LOB (Garside *et al.*, 1987), Susanne and SEC (Taylor and Knowles, 1988) corpora. We calculated *type precision* (percentage of SCFs acquired which were also exemplified in the manual analysis) and *recall* (percentage of the SCFs exemplified in the manual analysis which were also acquired automatically), and combined them into a single measure of overall performance using the *F measure* (Manning and Schütze, 1999).

$$F = \frac{2 \cdot \text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}} \quad (4)$$

We estimated *accuracy* with which the system ranks true positive classes against the correct ranking. This was computed by calculating the percentage of pairs of SCFs at positions (n, m) such that $n < m$ in the system ranking that occur in the same order in the ranking from the manual analysis. This gives us an estimate of the accuracy of the relative frequencies of SCFs output by the system. In addition to the system results, we also calculated KL and RC between the acquired unfiltered SCF distributions and the distributions obtained from the manual analysis.

4.2 Results

Table 5 gives average results for the 60 test verbs using each method. The results indicate that both add one smoothing and Katz backing-off improve the baseline performance only slightly. Linear interpolation outperforms these methods, achieving better results on all measures. The improved KL indicates that the method improves the overall accuracy of SCF distributions. The results with RC and system accuracy show that it helps to correct the ranking of SCFs. The fact

Method	KL	RC	System results (%)			
			accuracy	precision	recall	F measure
Baseline	0.63	0.72	79.2	78.5	63.3	70.1
Add one	0.64	0.74	79.0	85.3	59.7	70.2
Katz backing-off	0.61	0.75	79.0	76.4	67.6	71.7
Linear interpolation	0.51	0.82	84.4	87.8	68.7	77.1

Table 5: Average results with different methods using semantically motivated back-off estimates for smoothing

Method	KL	RC	System results (%)			
			accuracy	precision	recall	F measure
Baseline	0.63	0.72	79.2	78.5	63.3	70.1
Katz backing-off	0.68	0.69	77.2	75.2	61.7	67.8
Linear interpolation	0.79	0.64	76.7	71.4	64.1	67.6

Table 6: Average results using the SCF distribution of all verbs as back-off estimates for smoothing

that both precision and recall show clear improvement over the baseline results demonstrates that linear interpolation can be successfully combined with the filtering method employed. These results seem to suggest that a smoothing method which affects both the highly ranked SCFs and SCFs of low frequency is profitable for this task.

In this experiment, the semantically motivated back-off estimates helped to reduce the sparse data problem significantly. While a total of 151 correct SCFs were missing in the test data, only three were missing after smoothing with Katz backing-off or linear interpolation.

For comparison, we re-run these experiments using the general SCF distribution of all verbs as back-off estimates for smoothing⁶. The average results for the 60 test verbs given in table 6 show that when using these estimates, we obtain worse results than with the baseline method. This demonstrates that while such estimates provide an easy solution to the sparse data problem, they can actually degrade the accuracy of verbal acquisition.

Table 7 displays individual results for the different verb classes. It lists the results obtained with KL and RC using the baseline method and linear interpolation with semantically motivated estimates. Examining the results obtained with linear interpolation allows us to consider the accuracy of the back-off es-

timates for each verb class. Out of ten verb classes, eight show improvement with linear interpolation, with both KL and RC. However, two verb classes - aspectual verbs, and verbs of appearance, disappearance and occurrence - show worse results when linear interpolation is used.

According to Levin (1993), these two verb classes need further classification before a full semantic account can be made. The problem with aspectual verbs is that the class contains verbs taking sentential complements. As Levin does not classify verbs on basis of their sentential complement-taking properties, more classification work is required before we can obtain accurate SCF estimates for this type of verb.

The problem with verbs of appearance is more specific to the verb class. Levin remarks that the definition of appearance verbs may be too loose. In addition, there are significant syntactic differences between the verbs belonging to the different sub-classes.

This suggests that we should examine the degree of SCF correlation between verbs from different sub-classes before deciding on the final (sub-)class for which we obtain the estimates. As the results with the combined Levin classes show, estimates can also be successfully built using verbs from different Levin classes, provided that the classes are similar enough.

⁶These estimates were obtained by extracting the number of verbs which are members of each SCF class in the ANLT dictionary. See section 2 for details.

Verb class	KL		RC	
	baseline	linear i.	baseline	linear i.
putting	0.70	0.66	0.68	0.70
sending and carrying, exerting force	0.64	0.50	0.72	0.96
change of possession	0.61	0.60	0.61	0.75
assessment, searching	0.81	0.62	0.61	0.70
social interaction	0.65	0.58	0.72	0.80
killing	0.69	0.67	0.91	0.95
destroy	0.95	0.20	0.70	0.97
appearance, disappearance and occurrence	0.14	0.17	0.91	0.83
motion	0.66	0.58	0.56	0.66
aspectual	0.48	0.54	0.86	0.89

Table 7: Baseline and linear interpolation results for the verb classes

5 Conclusion

In this paper, we have shown that the verb form specific SCF distributions of semantically similar verbs correlate well. On the basis of this observation, we have proposed using verb class specific back-off estimates in SCF acquisition. Employing the SCF acquisition framework of Briscoe and Carroll (1997), we have demonstrated that these estimates can be used to improve SCF acquisition significantly, when combined with smoothing and a simple filtering method.

We have not yet explored the possibility of using the semantically motivated estimates with statistical filtering. In principle, this should help to improve the performance of the statistical methods which make use of back-off estimates. If filtering based on relative frequencies still achieves better results, it would be worth investigating ways of handling the low frequency data for integration with this method. As Korhonen, Gorrell and McCarthy (2000) discuss, any statistical filtering method would work better at low frequencies than the one applied, since this simply disregards all low frequency SCFs.

In addition to refining the filtering method, our future work will focus on integrating this approach with large-scale SCF acquisition. This will involve (i) defining the set of semantic verb classes across the lexicon, (ii) obtaining back-off estimates for each verb class, and (iii) implementing a method capable of automatically classifying verbs to semantic classes. The latter can be done by linking the Word-

Net synonym sets with semantic classes, using a similar method to that employed by Dorr (1997). With the research reported, verbs were classified to semantic classes according to their most frequent sense. While this approach proved satisfactory, our future work will include investigating ways of addressing the problem of polysemy better.

The manual effort needed for obtaining the back-off estimates was quite high for this preliminary experiment. However, our recent investigation shows that the total number of semantic classes across the whole lexicon is unlikely to exceed 50. This is because many of the Levin classes have proved similar enough in terms of SCF distributions that they can be combined together. Therefore the additional effort required to carry out the proposed work seems justified, given the accuracy enhancement reported.

6 Acknowledgements

I thank Ted Briscoe and Diana McCarthy for useful comments on this paper.

References

- Boguraev, B., Briscoe, E., Carroll, J., Carter, D. and Grover, C. 1987. The derivation of a grammatically-indexed lexicon from the Longman Dictionary of Contemporary English. In *Proceedings of the 25th Annual Meeting of the Association for Computational Linguistics*, Stanford, CA. 193–200.
- Brent, M. 1991. Automatic acquisition of subcategorization frames from untagged text.

- In *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, Berkeley, CA. 209–214.
- Brent, M. 1993. From grammar to lexicon: unsupervised learning of lexical syntax. *Computational Linguistics* 19.3: 243–262.
- Briscoe, E. and Carroll, J. 1993. Generalised probabilistic LR parsing for unification-based grammars. *Computational Linguistics* 19.1: 25–60.
- Briscoe, E.J. and J. Carroll 1997. Automatic extraction of subcategorization from corpora. In *Proceedings of the 5th ACL Conf. on Applied Natural Language Processing*, Washington, DC. 356–363.
- Briscoe, T., Carroll, J. and Korhonen, A. 1997. *Automatic extraction of subcategorization frames from corpora - a framework and 3 experiments*. Sparkle WP5 Deliverable. Available in <http://www.ilc.pi.cnr.it/>.
- Carroll, G. and Rooth, M. 1998. Valence induction with a head-lexicalized PCFG. In *Proceedings of the 3rd Conference on Empirical Methods in Natural Language Processing*, Granada, Spain.
- Cover, Thomas, M. and Thomas, J.A. 1991. *Elements of Information Theory*. Wiley-Interscience, New York.
- Dorr, B. 1997. Large-scale dictionary construction for foreign language tutoring and interlingual machine translation. *Machine Translation* 12.4: 271–325.
- Gahl, S. 1998. Automatic extraction of subcorpora based on subcategorization frames from a part-of-speech tagged corpus. In *Proceedings of the COLING-ACL'98*, Montreal, Canada.
- Grishman, R., Macleod, C. and Meyers, A. 1994. Complex syntax: building a computational lexicon. In *Proceedings of the International Conference on Computational Linguistics, COLING-94*, Kyoto, Japan. 268–272.
- Garside, R., Leech, G. and Sampson, G. 1987. *The computational analysis of English: A corpus-based approach*. Longman, London.
- Katz, S. M. 1987. Estimation of probabilities from sparse data for the language model component of speech recogniser. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 35.3: 400–401.
- Korhonen, A., Gorrell, G. and McCarthy, D. 2000. Statistical filtering and subcategorization frame acquisition. In *Proceedings of the Joint SIGDAT Conference on Empirical Methods in Natural Language Processing and Very Large Corpora*, Hong Kong.
- Leech, G. 1992. 100 million words of English: the British National Corpus. *Language Research* 28(1): 1–13.
- Levin, B. 1993. *English Verb Classes and Alternations*. Chicago University Press, Chicago.
- Manning, C. 1993. Automatic acquisition of a large subcategorization dictionary from corpora. In *Proceedings of the 31st Annual Meeting of the Association for Computational Linguistics*, Columbus, Ohio. 235–242.
- Manning, C. and Schütze, H. 1999. *Foundations of Statistical Natural Language Processing*. New York University, Ms.
- Miller, G., Beckwith, R., Felbaum, C., Gross, D. and Miller, K. 1993. *Introduction to WordNet: An On-Line Lexical Database*. <ftp://clarity.princeton.edu/pub/WordNet/5papers.ps>.
- Sampson, G. 1995. *English for the computer*. Oxford, UK: Oxford University Press.
- Sarkar, A. and Zeman, D. 2000. Automatic Extraction of Subcategorization Frames for Czech. In *Proceedings of the International Conference on Computational Linguistics, COLING-00*, Saarbrücken, Germany. 691–697.
- Spearman, C. 1904. The proof and measurement of association between two things. *American Journal of Psychology* 15: 72–101.
- Taylor, L. and Knowles, G. 1988. *Manual of information to accompany the SEC corpus: the machine-readable corpus of spoken English*. University of Lancaster, UK, Ms.
- Ushioda, A., Evans, D., Gibson, T. and Waibel, A. 1993. The automatic acquisition of frequencies of verb subcategorization frames from tagged corpora. In Boguraev, B. and Pustejovsky, J. eds. *SIGLEX ACL Workshop on the Acquisition of Lexical Knowledge from Text*. Columbus, Ohio: 95–106.