

A Comparison between Supervised Learning Algorithms for Word Sense Disambiguation

Gerard Escudero, Lluís Màrquez, and German Rigau*

TALP Research Center. LSI Department. Universitat Politècnica de Catalunya (UPC)

Jordi Girona Salgado 1–3. E-08034 Barcelona. Catalonia

{escudero,lluism,g.rigau}@lsi.upc.es

Abstract

This paper describes a set of comparative experiments, including cross-corpus evaluation, between five alternative algorithms for supervised Word Sense Disambiguation (WSD), namely Naive Bayes, Exemplar-based learning, *SNoW*, Decision Lists, and Boosting. Two main conclusions can be drawn: 1) The *LazyBoosting* algorithm outperforms the other four state-of-the-art algorithms in terms of accuracy and ability to tune to new domains; 2) The domain dependence of WSD systems seems very strong and suggests that some kind of adaptation or tuning is required for cross-corpus application.

1 Introduction

Word Sense Disambiguation (WSD) is the problem of assigning the appropriate meaning (or sense) to a given word in a text or discourse. Resolving the ambiguity of words is a central problem for large scale language understanding applications and their associate tasks (Ide and Véronis, 1998). Besides, WSD is one of the most important open problems in NLP. Despite the wide range of approaches investigated (Kilgarriff and Rosenzweig, 2000) and the large effort devoted to tackle this problem, to date, no large-scale broad-coverage and highly accurate WSD system has been built.

One of the most successful current lines of research is the corpus-based approach in which statistical or Machine Learning (ML) algorithms have been applied to learn statistical models or classifiers from corpora in order to perform

WSD. Generally, supervised approaches (those that learn from previously semantically annotated corpus) have obtained better results than unsupervised methods on small sets of selected ambiguous words, or artificial pseudo-words. Many standard ML algorithms for supervised learning have been applied, such as: Decision Lists (Yarowsky, 1994; Agirre and Martínez, 2000), Neural Networks (Towell and Voorhees, 1998), Bayesian learning (Bruce and Wiebe, 1999), Exemplar-Based learning (Ng, 1997a), and Boosting (Escudero et al., 2000a), etc. Further, in (Mooney, 1996) some of the previous methods are compared jointly with Decision Trees and Rule Induction algorithms, on a very restricted domain.

Although some comparative studies between alternative algorithms have been reported (Mooney, 1996; Ng, 1997a; Escudero et al., 2000a; Escudero et al., 2000b), none of them addresses the issue of the portability of supervised ML algorithms for WSD, i.e. to test whether the accuracy of a system trained on a certain corpus can be extrapolated to other corpora or not. We think that the study of the domain dependence of WSD—in the style of other studies devoted to parsing (Sekine, 1997; Ratnaparkhi, 1999)—is needed to assess the validity of the supervised approach, and to determine to which extent a pre-process of tuning is necessary to make real WSD systems portable. In this direction, this work compares five different ML algorithms and explores their portability and tuning ability by training and testing them on different corpora.

2 Learning Algorithms Tested

Naive-Bayes (NB). Naive Bayes is intended as a simple representative of statistical learning methods. It has been used in its most classical

* This research has been partially funded by the Spanish Research Department (CICYT's project TIC98-0423-C06), by the EU Commission (NAMIC IST-1999-12392), and by the Catalan Research Department (CIRIT's consolidated research group 1999SGR-150 and CIRIT's grant 1999FI 00773).

setting (Duda and Hart, 1973). That is, assuming independence of features, it classifies a new example by assigning the class that maximizes the conditional probability of the class given the observed sequence of features of that example.

Model probabilities are estimated during training process using relative frequencies. To avoid the effect of zero counts when estimating probabilities, a very simple smoothing technique has been used, which was proposed in (Ng, 1997a).

Despite its simplicity, Naive Bayes is claimed to obtain state-of-the-art accuracy on supervised WSD in many papers (Mooney, 1996; Ng, 1997a; Leacock et al., 1998).

Exemplar-based Classifier (EB). In exemplar, instance, or memory-based learning (Aha et al., 1991) no generalization of training examples is performed. Instead, the examples are stored in memory and the classification of new examples is based on the classes of the most similar stored examples. In our implementation, all examples are kept in memory and the classification of a new example is based on a k -NN (Nearest-Neighbours) algorithm using Hamming distance to measure closeness. For k 's greater than 1, the resulting sense is the weighted majority sense of the k nearest neighbours —where each example votes its sense with a strength proportional to its closeness to the test example.

Exemplar-based learning is said to be the best option for WSD (Ng, 1997a). Other authors (Daelemans et al., 1999) point out that exemplar-based methods tend to be superior in language learning problems because they do not forget exceptions.

SNoW: A Winnow-based Classifier. *SNoW* stands for Sparse Network Of Winnows, and it is intended as a representative of on-line learning algorithms. In the *SNoW* architecture there is a Winnow (Littlestone, 1988) node for each class, which learns to separate that class from all the rest. In this paper, our approach to WSD using *SNoW* follows that of (Escudero et al., 2000c).

SNoW is proven to perform very well in high dimensional domains, where both, the training examples and the target function reside very sparsely in the feature space (Roth, 1998), e.g: text categorization, context-sensitive spelling

correction, WSD, etc.

Decision Lists (DL). In this setting, Decision Lists are ordered lists of features extracted from the training examples and weighted by a log-likelihood measure (Yarowsky, 1994). The approximation described in (Agirre and Martínez, 2000) has been fully used (using also their pruning and smoothing techniques).

Decision Lists were one of the most successful systems on the 1st edition of the Senseval competition (Kilgarriff and Rosenzweig, 2000).

LazyBoosting (LB). The main idea of boosting algorithms is to combine many simple and moderately accurate hypotheses (called weak classifiers) into a single, highly accurate classifier. The weak classifiers are trained sequentially and, conceptually, each of them is trained on the examples which were most difficult to classify by the preceding weak classifiers.

LazyBoosting (Escudero et al., 2000a), is a simple modification of the AdaBoost.MH algorithm (Schapire and Singer, to appear), which consists of reducing the feature space that is explored when learning each weak classifier. More specifically, a small proportion of attributes are randomly selected and the best weak rule is selected only among them. This modification significantly increases the efficiency of the learning process with no loss in accuracy.

3 Setting

The set of comparative experiments has been carried out on a subset of 21 words of the DSO corpus, which is a semantically annotated English corpus collected by Ng and colleagues (Ng and Lee, 1996), and available from the Linguistic Data Consortium (LDC)¹. Each word is treated as a different classification problem. They are 13 nouns (age, art, body, car, child, cost, head, interest, line, point, state, thing, work) and 8 verbs (become, fall, grow, lose, set, speak, strike, tell). The average number of senses per word is close to 10 and the number of training examples is close to 1,000.

The DSO corpus contains sentences from two different corpora, namely Wall Street Journal (WSJ) and Brown Corpus (BC). Therefore, it is easy to perform experiments about the portability of alternative systems by training them on

¹<http://www.ldc.upenn.edu/>

the WSJ part (A part, hereinafter) and testing them on the BC part (B part, hereinafter), or vice-versa.

Two kinds of information are used to train classifiers: local and topical context. The former consists of the words and part-of-speech tags appearing in a window of ± 3 items around the target word, and collocations of up to three consecutive words in the same window. The latter consists of the unordered set of content words appearing in the whole sentence.

4 Experiments

4.1 Comparing the five approaches

The five algorithms, jointly with a naive Most-Frequent-sense Classifier (MFC), have been tested on 7 different combinations of training–test sets². Accuracy figures, averaged over the 21 words, are reported in table 1. The comparison leads to the following conclusions:

LazyBoosting outperforms the other three methods in all tests. The difference is statistically significant in all cases except when comparing **LazyBoosting** to the Exemplar Based approach in the case marked with an asterisk³.

Extremely poor results are observed when testing the portability of the systems. Restricting to **LazyBoosting** results, we observe that the accuracy obtained in A–B is 47.1% while the accuracy in B–B (which can be considered an upper bound for **LazyBoosting** in B corpus) is 59.0%, that is, a drop of 12 points. Furthermore, 47.1% is only slightly better than the most frequent sense in corpus B, 45.5%.

Apart from accuracy figures, the observation of the predictions made by the five methods on the test sets provides interesting information about the comparison of the algorithms. Table 2 shows the agreement rates and the Kappa

(κ) statistics⁴ between all pairs of methods in the A+B–A+B case. ‘DSO’ stands for the annotation of DSO corpus, which is taken as the correct. Therefore the agreement rate with DSO contains the accuracy results previously reported. Some interesting conclusions can be drawn from those tables:

1. NB obtains the most similar results with regard to MFC in agreement rate and Kappa values in all tables. The agreement ratio is 76%, that is, more than 3 out of 4 times it predicts the most frequent sense.

2. LB obtains the most similar results with regard to DSO (accuracy) in agreement rate and Kappa values, and it has the less similar Kappa and agreement values with regard to MFC. This indicates that LB is the method that better learns the behaviour of the DSO examples.

3. The Kappa values are very low. But, as it is suggested in (Véronis, 1998), evaluation measures, such as precision and recall, should be computed relative to the agreement between the human annotators of the corpus and not to a theoretical 100%. It seems pointless to expect more agreement between the system and the reference corpus than between the annotators themselves. Contrary to the intuition that the agreement between human annotators should be very high in the WSD task, some papers report surprisingly low figures. For instance, (Ng et al., 1999) reports an accuracy rate of 56.7% and a Kappa value of 0.317 when comparing the annotation of a subset of the DSO corpus performed by two independent research groups⁵. Similarly, (Véronis, 1998) reports values of Kappa near to zero when annotating some special words for the ROMANSEVAL corpus⁶. From this point of view, the Kappa values of 0.44 achieved by LB in A+B–A+B could be considered excellent results. Unfortunately, the subset of the DSO corpus and that used in this

²The combinations of training–test sets are called: A+B–A+B, A+B–A, A+B–B, A–A, B–B, A–B, and B–A, respectively. In this notation, the training set is placed at the left hand side of symbol “–”, while the test set is at the right hand side. For instance, A–B means that the training set is corpus A and the test set is corpus B. The symbol “+” stands for set union.

³Statistical tests of significance applied: McNemar’s test and 10-fold cross-validation paired Student’s *t*-test at a confidence value of 95% (Dietterich, 1998).

⁴The Kappa statistic k (Cohen, 1960) is a better measure of inter-annotator agreement which reduces the effect of chance agreement. It has been used for measuring inter-annotator agreement during the construction of some semantic annotated corpora (Véronis, 1998; Ng et al., 1999).

⁵A Kappa value of 1 indicates perfect agreement, while 0.8 is considered as indicating good agreement (Carletta, 1996).

⁶<http://www.lpl.univ-aix.fr/projects/romanseval>

	Accuracy (%)						
	A+B-A+B	A+B-A	A+B-B	A-A	B-B	A-B	B-A
MFC	46.55±0.71	53.90±2.01	39.21±1.90	55.94±1.10	45.52±1.27	36.40	38.71
Naive Bayes	61.55±1.04	67.25±1.07	55.85±1.81	65.86±1.11	56.80±1.12	41.38	47.66
Exemplar-based	63.01±0.93	69.08±1.66	56.97±1.22	68.98±1.06	57.36±1.68	45.32	51.13
Decision Lists	61.58±0.98	67.64±0.94	55.53±1.85	67.57±1.44	56.56±1.59	43.01	48.83
<i>SNoW</i>	60.92±1.09	65.57±1.33	56.28±1.10	67.12±1.16	56.13±1.23	44.07	49.76
LazyBoosting	66.32±1.34	71.79±1.51	60.85±1.81	71.26±1.15	58.96±1.86	47.10	51.99*

Table 1: Accuracy results (\pm standard deviation) of the methods on all training–test combinations

report are not the same and, therefore, a direct comparison is not possible.

	A+B-A+B						
	DSO	MFC	NB	EB	SN	DL	LB
DSO	—	46.6	61.6	63.0	60.9	61.6	66.3
MFC	-0.19	—	73.9	60.0	55.9	64.9	54.9
NB	0.24	-0.09	—	76.3	74.5	76.8	71.4
EB	0.36	-0.15	0.44	—	69.6	70.7	72.5
SN	0.36	-0.17	0.44	0.44	—	67.5	69.0
DL	0.32	-0.13	0.40	0.41	0.38	—	69.9
LB	0.44	-0.17	0.37	0.50	0.46	0.42	—

Table 2: Kappa (κ) statistic (below diagonal) and agreement rate (above diagonal) between all methods in A+B-A+B experiments

4.2 About the tuning to new domains

This experiment explores the effect of a simple tuning process consisting of adding to the original training set a relatively small sample of manually sense tagged examples of the new domain. The size of this supervised portion varies from 10% to 50% of the available corpus in steps of 10% (the remaining 50% is kept for testing). Results indicate that: **LazyBoosting** is again superior to their competitors.

Summarizing, the results obtained show that for Naive Bayes, Exemplar Based, *SNoW* and Decision Lists methods it is not worth keeping the original training examples. Instead, a better (but disappointing) strategy would be simply using the tuning corpus. However, this is not the situation of **LazyBoosting**, for which a moderate (but consistent) improvement of accuracy is observed when retaining the original training set.

We observed that part of the poor results obtained is explained by: 1) Corpus A and B have a very different distribution of senses, and, therefore, different a-priori biases; Furthermore, 2) Examples of corpus A and B con-

tain different information, and, therefore, the learning algorithms acquire different (and non interchangeable) classification cues from both corpora. The study of the rules acquired by **LazyBoosting** from *WSJ* and *BC* helped understanding the differences between corpora. On the one hand, the type of features used in the rules were significantly different between corpora, and, additionally, there were very few rules that apply to both sets; On the other hand, the sign of the prediction of many of these common rules was somewhat contradictory between corpora (Escudero et al., 2000c).

4.3 About the training data quality

The observation of the rules acquired by **LazyBoosting** also could help improving data quality. It is known that mislabelled examples resulting from annotation errors tend to be hard examples to classify correctly, and, therefore, tend to have large weights in the final distribution. This observation allows both to identify the noisy examples and use **LazyBoosting** as a way to improve the training corpus.

A preliminary experiment has been carried out in this direction by studying the rules acquired by **LazyBoosting** from the training examples of word *state*. The manually revision of the 50 highest scored rules allowed us to identify a high number of noisy training examples –there were 11 of 50 tagging errors–, and, additionally, 17 examples of 50 not coherently tagged, probably due to the too fine grained or not so clear distinctions between the senses involved in these examples. Thus, there were 28 of 50 examples with some problem, that is more than 1 of each two cases have a problem.

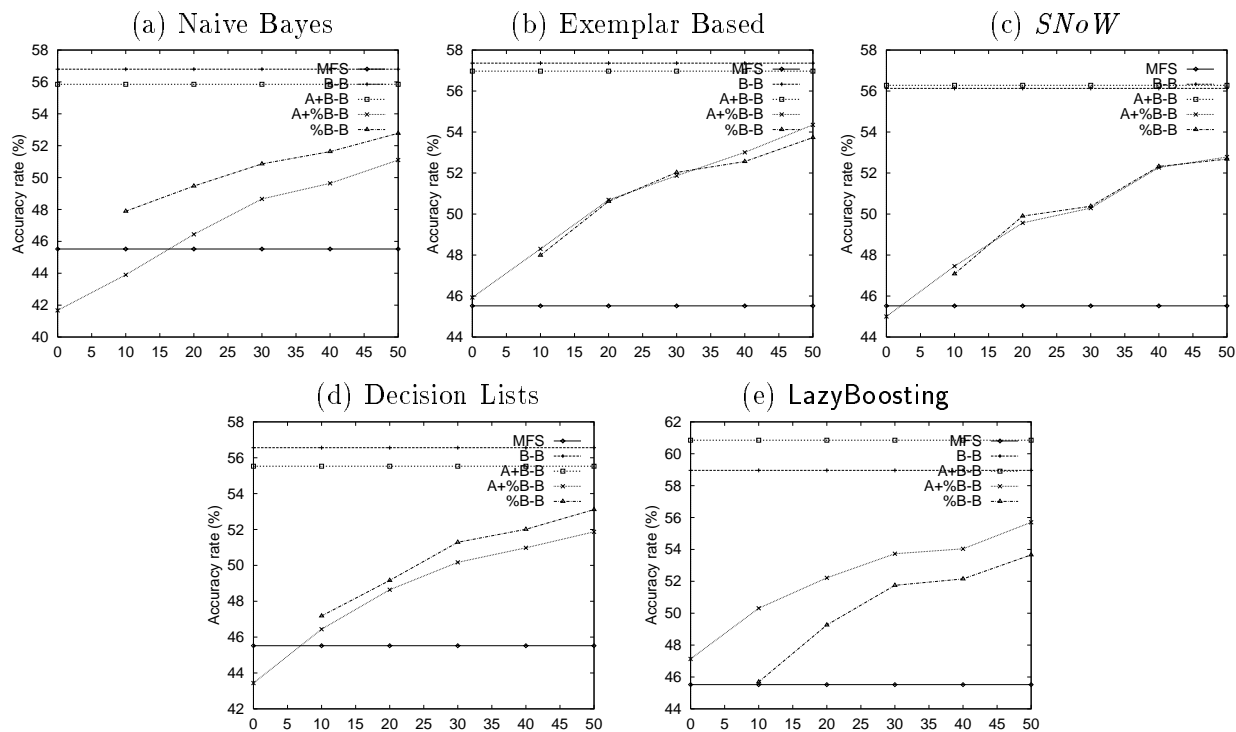


Figure 1: Results of the tuning experiment

5 Conclusions and Future Work

This work reports a comparative study of five ML approaches to WSD, and focuses on studying their portability. The main conclusions are:

LazyBoosting algorithm outperforms the other four state-of-the-art supervised ML methods in all domains tested. Furthermore, this algorithm shows better properties when is tuned to new domains.

Portability is a very important issue that has been paid little attention up to the present. In this paper we show that a process of tuning to the domain of application is required to assure the portability of WSD systems (at least if the learning-testing corpora differ so as BC and WSJ do). This evidence questions the idea of “robust broad-coverage WSD” introduced by (Ng, 1997b), in which a supervised system trained on a large enough corpora (say a thousand examples per word) should provide fairly accurate disambiguation on any corpora. To determine the viability of the supervised approach to WSD we believe that a serious effort should be devoted to study the problem of obtaining representative enough training corpora at a reasonable cost.

Further work is planned to be done in the following directions:

1. Since most of the knowledge learned from a domain is not useful when changing to a new domain, further investigation is needed on tuning strategies, specially on those using non-supervised algorithms.
2. It has been noted that mislabelled examples resulting from annotation errors tend to be hard examples to classify correctly, and, therefore, tend to have large weights in the final distribution. It could provide the methodologies to automatic verify the semantic annotation of corpora and the grouping of senses.
3. Moreover, the inspection of the rules learned by **LazyBoosting** could provide evidence about similar behaviours of a-priori different senses. This type of knowledge could be useful to perform clustering of too fine-grained or artificial senses.

References

- E. Agirre and D. Martínez. 2000. Decision Lists and Automatic Word Sense Disambiguation. In *Proceedings of the COLING Workshop on Seman-*

- tic Annotation and Intelligent Content*, Luxembourg.
- D. W. Aha, D. Kibler, and M. K. Albert. 1991. Instance-based learning algorithms. *Machine Learning*, 6(1):37–66.
- R. F. Bruce and J. M. Wiebe. 1999. Decomposable Modeling in Natural Language Processing. *Computational Linguistics*, 25(2):195–207.
- J. Carletta. 1996. Assessing agreement on classification tasks: The kappa statistic. *Computational Linguistics*, 22(2):249–254.
- J. Cohen. 1960. A Coefficient of Agreement for Nominal Scales. *Journal of Educational and Psychological Measurement*, 20:37–46.
- W. Daelemans, A. van den Bosch, and J. Zavrel. 1999. Forgetting Exceptions is Harmful in Language Learning. *Machine Learning*, 34:11–41.
- T. G. Dietterich. 1998. Approximate Statistical Tests for Comparing Supervised Classification Learning Algorithms. *Neural Computation*, 10(7).
- R. O. Duda and P. E. Hart. 1973. *Pattern Classification and Scene Analysis*. John Wiley & Sons, New York.
- G. Escudero, L. Màrquez, and G. Rigau. 2000a. Boosting Applied to Word Sense Disambiguation. In *Proceedings of the 12th European Conference on Machine Learning, ECML*, Barcelona, Spain.
- G. Escudero, L. Màrquez, and G. Rigau. 2000b. Naive Bayes and Exemplar-Based Approaches to Word Sense Disambiguation Revisited. In *Proceedings of the 14th European Conference on Artificial Intelligence, ECAI*, Berlin, Germany.
- G. Escudero, L. Màrquez, and G. Rigau. 2000c. On the Portability and Tuning of Supervised Word Sense Disambiguation Systems. Research Report LSI-00-30-R, Software Department (LSI). Technical University of Catalonia (UPC), Barcelona.
- A. Fujii, K. Inui, T. Tokunaga, and H. Tanaka. 1998. Selective Sampling for Example-Based Word Sense Disambiguation. *Computational Linguistics*, 24(4):573–597.
- A. R. Golding and D. Roth. 1999. A Winnow-Based Approach to Context-Sensitive Spelling Correction. *Machine Learning*, 34:107–130.
- N. Ide and J. Véronis. 1998. Introduction to the Special Issue on Word Sense Disambiguation: The State of the Art. *Computational Linguistics*, 24(1):1–40.
- A. Kilgarriff and J. Rosenzweig. 2000. English SENSEVAL: Report and Results. In *Proceedings of the 2nd International Conference on Language Resources and Evaluation, LREC*, Athens, Greece.
- C. Leacock, M. Chodorow, and G. A. Miller. 1998. Using Corpus Statistics and WordNet Relations for Sense Identification. *Computational Linguistics*, 24(1):147–166.
- N. Littlestone. 1988. Learning Quickly when Irrelevant Attributes Abound. *Machine Learning*, 2:285–318.
- R. J. Mooney. 1996. Comparative Experiments on Disambiguating Word Senses: An Illustration of the Role of Bias in Machine Learning. In *Proceedings of the 1st Conference on Empirical Methods in Natural Language Processing, EMNLP*.
- H. T. Ng and H. B. Lee. 1996. Integrating Multiple Knowledge Sources to Disambiguate Word Sense: An Exemplar-based Approach. In *Proceedings of the 34th Annual Meeting of the Association for Computational Linguistics, ACL*.
- H. T. Ng, C. Y. Lim, and S. K. Foo. 1999. A Case Study on Inter-Annotator Agreement for Word Sense Disambiguation. In *Proceedings of the ACL SIGLEX Workshop: Standardizing Lexical Resources*.
- H. T. Ng. 1997a. Exemplar-Base Word Sense Disambiguation: Some Recent Improvements. In *Proceedings of the 2nd Conference on Empirical Methods in Natural Language Processing, EMNLP*.
- H. T. Ng. 1997b. Getting Serious about Word Sense Disambiguation. In *Proceedings of the ACL SIGLEX Workshop: Standardizing Lexical Resources*.
- A. Ratnaparkhi. 1999. Learning to Parse Natural Language with Maximum Entropy Models. *Machine Learning*, 34:151–175.
- D. Roth. 1998. Learning to Resolve Natural Language Ambiguities: A Unified Approach. In *Proceedings of the 15th Conference of the American Association for Artificial Intelligence, AAAI*, pages 806–813.
- R. E. Schapire and Y. Singer. to appear. Improved Boosting Algorithms Using Confidence-rated Predictions. *Machine Learning*.
- S. Sekine. 1997. The Domain Dependence of Parsing. In *Proceedings of the 5th ACL Conference on Applied Natural Language Processing, ANLP*, Washington DC, US.
- G. Towell and E. M. Voorhees. 1998. Disambiguating Highly Ambiguous Words. *Computational Linguistics*, 24(1):125–146.
- J. Véronis. 1998. A study of polysemy judgements and inter-annotator agreement. In *Programme and advanced papers of the Senseval workshop*, Herstmonceux Castle, England.
- D. Yarowsky. 1994. Decision Lists for Lexical Ambiguity Resolution: Application to Accent Restoration in Spanish and French. In *Proceedings of the 32nd Annual Meeting of the Association for Computational Linguistics, ACL*, pages 88–95, Las Cruces, NM, US.