



University of Pennsylvania
ScholarlyCommons

Departmental Papers (CIS)

Department of Computer & Information Science

March 2004

Flexible Margin Selection for Reranking with Full Pairwise Samples

Libin Shen
University of Pennsylvania

Aravind K. Joshi
University of Pennsylvania, joshi@cis.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/cis_papers

Recommended Citation

Libin Shen and Aravind K. Joshi, "Flexible Margin Selection for Reranking with Full Pairwise Samples", . March 2004.

Postprint version. Published in *Lecture Notes in Computer Science*, Volume 3248, Natural Language Processing - IJCNLP 2004: First International Joint Conference, 2005, pages 446-455.
Publisher URL: <http://dx.doi.org/10.1007/b105612>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/cis_papers/227
For more information, please contact repository@pobox.upenn.edu.

Flexible Margin Selection for Reranking with Full Pairwise Samples

Abstract

Perceptron like large margin algorithms are introduced for the experiments with various margin selections. Compared to the previous perceptron reranking algorithms, the new algorithms use full pairwise samples and allow us to search for margins in a larger space. Our experimental results on the data set of [1] show that a perceptron like ordinal regression algorithm with uneven margins can achieve Recall/Precision of 89.5/90.0 on section 23 of Penn Treebank. Our result on margin selection can be employed in other large margin machine learning algorithms as well as in other NLP tasks.

Comments

Postprint version. Published in *Lecture Notes in Computer Science*, Volume 3248, Natural Language Processing - IJCNLP 2004: First International Joint Conference, 2005, pages 446-455.
Publisher URL: <http://dx.doi.org/10.1007/b105612>

Flexible Margin Selection for Reranking with Full Pairwise Samples

Libin Shen and Aravind K. Joshi

Department of Computer and Information Science
University of Pennsylvania
Philadelphia, PA 19104, U.S.A.
{libin, joshi}@linc.cis.upenn.edu

Abstract. Perceptron like large margin algorithms are introduced for the experiments with various margin selections. Compared to the previous perceptron reranking algorithms, the new algorithms use full pairwise samples and allow us to search for margins in a larger space. Our experimental results on the data set of [1] show that a perceptron like ordinal regression algorithm with uneven margins can achieve Recall/Precision of 89.5/90.0 on section 23 of Penn Treebank. Our result on margin selection can be employed in other large margin machine learning algorithms as well as in other NLP tasks.

1 Introduction

In recent years, the so-called *reranking* techniques [1] have been successfully used in parameter estimation in many applications which were previously modeled as generative models. A baseline generative model generates N-best candidates, and then these candidates are reranked by using a rich set of local and global features. Various machine learning algorithms have been adapted to the reranking tasks.

In the field of machine learning, a class of tasks, which are called *ranking* or *ordinal regression*, are similar to the reranking tasks in NLP. A primary motivation of the present paper is to apply ranking or ordinal regression algorithms to the reranking tasks in NLP, especially because we observe that there is no direct way to apply these ranking algorithms to reranking. More specifically, we will compare the existing reranking and ranking algorithms in the framework of *margin selection*. The goal then is to search for a desirable margin for the reranking tasks in NLP.

In order to experiment with various margins, we will introduce variants of the traditional perceptron algorithm [2, 3] for reranking, which allows the use of various margins; The training is also very fast. The basic idea of these perceptron like algorithms is that we dynamically search for pairs of inconsistent objects and use them to update the weight vector. Since the ranks are ordered, the dynamical search can be implemented efficiently.

Compared to previous work on perceptron for parse reranking [4], our new algorithms use full pairwise samples instead of partial pairwise samples. This allows us to search for margins desirable for reranking tasks in a larger space, which is unavailable in the previous work.

In this paper, we focus on the parse reranking task. However, the methods can, of course, be applied to other NLP reranking tasks. Our experimental results on the data set in [1] show that a perceptron like ordinal regression algorithm with uneven margins can achieve Recall/Precision of 89.5/90.0 on section 23 of WSJ PTB, which is comparable to 89.6/89.9 with the boosting algorithm in [1], although boosting is believed to have more generalization capability. Our results also show that the new margins introduced in this paper are superior to the margins used in the previous works on reranking. The results on margin selection can be employed in reranking systems based on other machine learning algorithms, such as Winnow, Boosting and SVMs, as well as other NLP tasks, e.g. machine translation reranking.

2 Previous Works

2.1 Reranking

In recent years, reranking has been successfully applied to some NLP problems, especially to the problem of parse reranking. Ratnaparkhi [5] noticed that by ranking the 20-best parsing results generated by his maximal entropy parser, the F-measure went to 93% from 87%, if the oracle parse was successfully detected. Charniak [6] reranked the N-best parses by reestimating a language model on a large number of features.

Collins [1] first used machine learning algorithms for parse reranking. Two approaches were proposed in that paper; one used Boosting Loss and the other used Log-Likelihood Loss. Boosting Loss achieved better results. The Boosting Loss model is as follows. Let $\mathbf{x}_{i,j}$ be the feature vector of the j^{th} parse of the i^{th} sentence. Let $\tilde{\mathbf{x}}_i$ be the feature vector of the best parse for the i^{th} sentence. Let F_α be a score function

$$F_\alpha(\mathbf{x}_{i,j}) \equiv \alpha' \cdot \mathbf{x}_{i,j},$$

where α is a weight vector. The *margin* $M_{\alpha,i,j}$ on sample $\mathbf{x}_{i,j}$ is defined as

$$M_{\alpha,i,j} \equiv F_\alpha(\tilde{\mathbf{x}}_i) - F_\alpha(\mathbf{x}_{i,j})$$

Finally the Boost Loss function is defined as

$$BoostLoss(\alpha) \equiv \sum_i \sum_j e^{-(F_\alpha(\tilde{\mathbf{x}}_i) - F_\alpha(\mathbf{x}_{i,j}))} = \sum_i \sum_j e^{-M_{\alpha,i,j}}$$

The Boosting algorithm was used to search the weight vector α to minimize the Boost Loss.

We may rewrite the definition of the margin $M_{\alpha,i,j}$ by using pairwise samples as follows.

$$\mathbf{s}_{i,j} \equiv \tilde{\mathbf{x}}_i - \mathbf{x}_{i,j}, \text{ then}$$

$$M_{\alpha,i,j} = F_\alpha(\tilde{\mathbf{x}}_i) - F_\alpha(\mathbf{x}_{i,j}) = F_\alpha(\tilde{\mathbf{x}}_i - \mathbf{x}_{i,j}) = F_\alpha(\mathbf{s}_{i,j})$$

So the Boosting Loss approach in [1] is similar to maximizing the margin [7] between 0 and $F_\alpha(\mathbf{s}_{i,j})$, where $\mathbf{s}_{i,j}$ are pairwise samples as we have described above.

In [4], the voted perceptron and the Tree kernel were applied to parse reranking. Similar to [1], pairwise samples were used as training samples. The perceptron updating step was defined as

$$\mathbf{w}^{t+1} = \mathbf{w}^t + \tilde{\mathbf{x}}_i - \mathbf{x}_{i,j},$$

where \mathbf{w}^t is the weight vector at the t th updating. This is equivalent to using pairwise sample $\mathbf{s}_{i,j}$ as we have defined above.

$$\mathbf{w}^{t+1} = \mathbf{w}^t + \mathbf{s}_{i,j}$$

In our previous work [8], we applied Support Vector Machines (SVMs) and Tree kernels to parse reranking. In that paper, pairwise samples were used explicitly through the Preference kernel. $\mathbf{u}_{i,j}^+$ and $\mathbf{u}_{i,j}^-$ defined as follows were used as positive samples and negative samples respectively.

$$\mathbf{u}_{i,j}^+ \equiv (\tilde{\mathbf{x}}_i, \mathbf{x}_{i,j}), \quad \mathbf{u}_{i,j}^- \equiv (\mathbf{x}_{i,j}, \tilde{\mathbf{x}}_i)$$

SVM is used to maximize the margin between positive samples and negative samples, which in turn is proportional to the margin between the best parse of each sentence and the rest of the N-best parses.

In the works on reranking, the margin is defined as the distance between the best candidate and the rest. The reranking problem is reduced to a classification problem by using pairwise samples implicitly or explicitly.

2.2 Ranking

In the previous works on ranking or ordinal regression, the margin is defined as the distance between two consecutive ranks. Two approaches have been used. One is *PRanking* that extends the perceptron algorithm by using multiple biases to represent the boundaries between every two consecutive ranks [9]. However, due to the introduction of a set of biases it is impossible to use PRanking in other ranking-like problems. The other approach is to reduce the ranking problem to a classification problem by using the trick of pairwise samples [10].

2.3 Large Margin Classifiers

There are quite a few linear classifiers¹ that can separate samples with large margin, such as SVMs [11], Boosting [7], Winnow [12] and Perceptron [13]. The performance of SVMs is superior to other linear classifiers because of their ability to maximize the margin, but SVMs are slow in training.

For margin selection, we do need an algorithm that runs fast for training, so that we can test various margins. Then the result of the margin selection can be employed in other linear classifiers. For the purpose of margin selection we propose perceptron like algorithms for the following two reasons. First, perceptron is fast in training which allows us to do experiments with various margin selections on real-world data. Furthermore, perceptron algorithms are simple in principle, which makes it easy to implement modification.

¹ Here we do not consider kernels of infinite dimension

3 Ranks and Margins for Reranking

In the previous works on ranking, ranks are defined on the whole training and test data. Thus we can define boundaries between consecutive ranks on the whole data. In the reranking problem, ranks are *local*. They are defined over a sub set of the samples in the data set. For example, in the parse reranking problem, the rank of a parse is only defined as the rank among all the parses for the same sentence. The training data includes 36,000 sentence, with an average of about 27 parses per sentence [1].

As a result, we cannot use the PRank algorithm in the reranking task, since there are no *global* ranks or boundaries in reranking, as the PRank algorithm is designed to estimate the global rank boundaries over all the samples during the training. If we introduce auxiliary variables for the boundaries for each cluster, the number of the parameters will be as large as the number of samples. Obviously this is not a good idea. However, the approach of using pairwise samples works. By pairing up two samples, we actually compute the relative distance between these two samples in the scoring metric.

Let r_i be the candidate parse that ranks as the i^{th} best for a sentence. The parses of the same sentence are ranked with respect to their *f-scores*, which measure the similarity to the gold standard parse. A parse with a large *f-score* is assigned a high rank. In reranking tasks, the margins between the best candidate and the rest are more useful. A hyperplane successfully separating r_1 and $r_2 \dots r_N$ is more predictive than a hyperplane successfully separating $r_1 \dots r_{10}$ and $r_{11} \dots r_N$, if we are only interested in the topmost result in test. This is also how the existing reranking systems are designed. However there are some problems with this approach.

There is a practical problem for the definition of the *best* parse in a sentence. In parse reranking, we may find several best parses for each training sentence instead of one. In order to break the tie, usually one selects just one of them arbitrarily as the top ranked parse and discard all others.

Furthermore, if we only look for the hyperplane to separate the best one from the rest, we, in fact, discard the order information of $r_2 \dots r_N$. For example, we did not employ the information that r_{10} is better than r_{11} in the training. Knowing r_{10} is better than r_{11} may be useless for training to some extent, but knowing r_2 is better than r_{11} is useful.

On the other the hand, the resulting weight vector w is supposed to assign the highest score to r_1 . Should it not assign the second highest score to r_2 ? Although we cannot give an affirmative answer at this time, it is at least reasonable to use more pairwise samples. This approach was avoided in the previous works on reranking, due to the problem of complexity of both the data size and the execution time. Thus we have provided a strong motivation for investigating some new reranking algorithms such that

- They utilize all the ordinal relations encoded in the ranked lists.
- The size of training data is the same as the original size of the ranked lists.
- The training time increases only moderately, although more information is used in training.

4 Perceptron for Ordinal Regression

4.1 Ordinal Regression

Let $\mathbf{x}_{i,j}, \mathbf{x}_{i,l}$ be the feature vectors of two parses for sentence i and $y_{i,j}, y_{i,l}$ be their ranks respectively, where $y_{i,j} + \epsilon < y_{i,l}$, and ϵ is a non-negative real number. It means that the rank of $\mathbf{x}_{i,j}$ is ϵ higher than the rank of $\mathbf{x}_{i,l}$. In this case, we say $\mathbf{x}_{i,j}$ is *significantly better* than $\mathbf{x}_{i,l}$. We are interested in finding a weight vector \mathbf{w} , such that

$$\mathbf{w} \cdot \mathbf{x}_{i,j} > \mathbf{w} \cdot \mathbf{x}_{i,l} + \tau, \text{ if } y_{i,j} + \epsilon < y_{i,l}$$

We ignore any pair of parses in which the difference in the ranks is $\leq \epsilon$. Hence, this problem is called ϵ -insensitive ordinal regression.

Let the training samples be

$$S = \{(\mathbf{x}_{i,j}, y_{i,j}) \mid 1 \leq i \leq m, 1 \leq j \leq k\},$$

where m is the number of sentences and k is the size of the N-best list. Let $f(\mathbf{x}) = \mathbf{w} \cdot \mathbf{x}$. We say the training data is ϵ -distinguishable by f if

$$\mathbf{w} \cdot \mathbf{x}_{i,j} > \mathbf{w} \cdot \mathbf{x}_{i,l}, \text{ if } y_{i,j} + \epsilon < y_{i,l},$$

for $1 \leq i \leq m, 1 \leq j, l \leq k$.

4.2 Dynamic Pairing

A straightforward method of using pairwise samples is to define positive and negative samples on the differences of vectors as in [10]. For each sentence i , $\mathbf{x}_{i,j} - \mathbf{x}_{i,l}$ is a positive sample if $y_{i,j} < y_{i,l}$, where $y_{i,j}$ is the rank of parse $\mathbf{x}_{i,j}$. Similarly, $\mathbf{x}_{i,j} - \mathbf{x}_{i,l}$ is a negative sample if $y_{i,l} < y_{i,j}$.

However, for real tasks, this greatly increases the data complexity from $O(mk)$ to $O(mk^2)$, where m is the number of training sentences, and k is the size of n-best list. For parse reranking k is about 27, and for machine translation reranking k is about 1000. Due to the limit of memory space we cannot define pairwise samples explicitly in this way.

The method to avoid this problem is to look up pairwise samples dynamically, as shown in **Algorithm 1**, a perceptron like algorithm. The basic idea is that, for each pair of parses for the same sentence, if

- the rank of $\mathbf{x}_{i,j}$ is significantly higher than the rank of $\mathbf{x}_{i,l}$, $y_{i,j} + \epsilon < y_{i,l}$
- the weight vector \mathbf{w} can not successfully separate $(\mathbf{x}_{i,j}$ and $\mathbf{x}_{i,l})$ with a learning margin τ , $\mathbf{w} \cdot \mathbf{x}_{i,j} < \mathbf{w} \cdot \mathbf{x}_{i,l} + \tau$,

then we need to update \mathbf{w} with the addition of $\mathbf{x}_{i,j} - \mathbf{x}_{i,l}$. It is not difficult to show Algorithm 1 is equivalent to using pairwise samples in training.

Algorithm 1 ordinal regression

Require: a positive learning margin τ .

```
1:  $t \leftarrow 0$ , initialize  $\mathbf{w}^0$ ;  
2: repeat  
3:   for (sentence  $i = 1, \dots, m$ ) do  
4:     for ( $1 \leq j < l \leq k$ ) do  
5:       if ( $y_{i,l} - y_{i,j} > \epsilon$  and  $\mathbf{w}^t \cdot \mathbf{x}_{i,j} < \mathbf{w}^t \cdot \mathbf{x}_{i,l} + \tau$ ) then  
6:          $\mathbf{w}^{t+1} \leftarrow \mathbf{w}^t + \mathbf{x}_{i,j} - \mathbf{x}_{i,l}$ ;  $t \leftarrow t + 1$ ;  
7:       else if ( $y_{i,j} - y_{i,l} > \epsilon$  and  $\mathbf{w}^t \cdot \mathbf{x}_{i,l} < \mathbf{w}^t \cdot \mathbf{x}_{i,j} + \tau$ ) then  
8:          $\mathbf{w}^{t+1} \leftarrow \mathbf{w}^t + \mathbf{x}_{i,l} - \mathbf{x}_{i,j}$ ;  $t \leftarrow t + 1$ ;  
9:       end if  
10:    end for  
11:  end for  
12: until no updates made in the outer for loop
```

Algorithm 2 ordinal regression, sentence updating

Require: a positive learning margin τ .

```
1:  $t \leftarrow 0$ , initialize  $\mathbf{w}^0$ ;  
2: repeat  
3:   for (sentence  $i = 1, \dots, m$ ) do  
4:     compute  $\mathbf{w}^t \cdot \mathbf{x}_{i,j}$  and  $u_j \leftarrow 0$  for all  $j$ ;  
5:     for ( $1 \leq j < l \leq k$ ) do  
6:       if ( $y_{i,l} - y_{i,j} > \epsilon$  and  $\mathbf{w}^t \cdot \mathbf{x}_{i,j} < \mathbf{w}^t \cdot \mathbf{x}_{i,l} + \tau$ ) then  
7:          $u_j \leftarrow u_j + 1$ ;  $u_l \leftarrow u_l - 1$ ;  
8:       else if ( $y_{i,j} - y_{i,l} > \epsilon$  and  $\mathbf{w}^t \cdot \mathbf{x}_{i,l} < \mathbf{w}^t \cdot \mathbf{x}_{i,j} + \tau$ ) then  
9:          $u_j \leftarrow u_j - 1$ ;  $u_l \leftarrow u_l + 1$ ;  
10:      end if  
11:    end for  
12:     $\mathbf{w}^{t+1} \leftarrow \mathbf{w}^t + \sum_j u_j \mathbf{x}_{i,j}$ ;  $t \leftarrow t + 1$ ;  
13:  end for  
14: until no updates made in the outer for loop
```

4.3 Sentence-Level Updating

In Algorithm 1, for each **repeat** iteration, the complexity is $O(mk^2d)$, where m and k are defined as above, and d is the average number of active features in a sample. We notice that the score of a parse $\mathbf{x}_{i,j}$ will be computed for k times in each **repeat** iteration. However, in many cases this is not necessary. In this section, we will revise Algorithm 1 to speed up the training phase.

Algorithm 2 is similar to Algorithm 1 except that the updating is not executed until all the inconsistent pairs for the same sentence are found. Therefore we only need to compute $\mathbf{w} \cdot \mathbf{x}_{i,j}$ for only once in each **repeat** iteration. So the complexity of each **repeat** iteration is $O(mk^2 + mkd)$.

4.4 Uneven Margins

For ϵ -insensitive ordinal regression, suppose $\epsilon = 10$ and our ordinal regression algorithm made two errors. One is on $(\mathbf{r}_1, \mathbf{r}_{11})$, and the other is on $(\mathbf{r}_{21}, \mathbf{r}_{31})$. The algorithm cannot recognize that the former is more serious than the latter. On the other hand, the algorithm does not try to distinguish \mathbf{r}_1 and \mathbf{r}_{10} , which is even worse.

Our solution is to apply *uneven margins* to the ϵ -insensitive ordinal regression. For example, we want to find a hyperplane for each sentence such that there is larger margin between \mathbf{r}_1 and \mathbf{r}_{10} , but a smaller margin between \mathbf{r}_1 and \mathbf{r}_2 , where \mathbf{r}_j is the parse that ranks j for a sentence. Similarly, we want a larger margin between \mathbf{r}_1 and \mathbf{r}_2 , but a smaller margin between \mathbf{r}_{10} and \mathbf{r}_{11} . Thus

$$\text{margin}(\mathbf{r}_1, \mathbf{r}_{10}) > \text{margin}(\mathbf{r}_1, \mathbf{r}_2) > \text{margin}(\mathbf{r}_{10}, \mathbf{r}_{11}) \quad (1)$$

So the solution is to search for a hyperplane such that

$$\text{score}(\mathbf{r}_p) - \text{score}(\mathbf{r}_q) > g(p, q)\tau$$

where $g(1, 10) > g(1, 2) > g(10, 11)$. Specifically, we replace one of the updating conditions

$$\mathbf{w} \cdot \mathbf{x}_{i,j} < \mathbf{w} \cdot \mathbf{x}_{i,l} + \tau$$

in line 6 of Algorithm 2 with

$$\frac{\mathbf{w} \cdot \mathbf{x}_{i,j} - \mathbf{w} \cdot \mathbf{x}_{i,l}}{g(y_{i,j}, y_{i,l})} < \tau, \quad (2)$$

and replace the updating condition

$$y_{i,l} - y_{i,j} > \epsilon \text{ with } g(y_{i,j}, y_{i,l}) > \epsilon, \quad (3)$$

which means that we ignore irrelevant inconsistent pairs with respect to g . We also replace the updating operation in line 7 with

$$u_j \leftarrow u_j + g(y_{i,j}, y_{i,l}), \quad u_l \leftarrow u_l - g(y_{i,j}, y_{i,l}) \quad (4)$$

A similar modification is made in line 8 and 9.

It can be shown that modifying Algorithm 2 in this way is equivalent to using $(\mathbf{x}_{i,j} - \mathbf{x}_{i,l})/g(y_{i,j}, y_{i,l})$ as pairwise samples, so it is well defined. Due to the space limitation, we omit the proof of the equivalence in this paper.

There are many candidates for the function g . The following function is one of the simplest solutions.

$$g(p, q) \equiv \frac{1}{p} - \frac{1}{q}$$

We will use this function in our experiments on parse reranking.

5 Experimental Results

In this section, we will report the experimental results for parse reranking task. We use the same data set as described in [1]. Section 2-21 of the WSJ Penn Treebank (PTB)[14] are used as training data, and section 23 is used for test. The training data contains around 40,000 sentences, each of which has 27 distinct parses on average. Of the 40,000 training sentences, the first 36,000 are used to train perceptrons. The remaining 4,000 sentences are used as development data for parameter estimation, such as the number of rounds of iteration in training. The 36,000 training sentences contain 1,065,620 parses totally. We use the feature set generated by Collins [1].

In all of our experiments, we have employed the voted perceptron as in [15, 4]. The voted version makes the result on the test set more stable.

In the first set of experiments, Algorithm 2 and its uneven margin variants are used. In addition, we evaluate the performance of separating only the best parse from the rest in training by modifying the updating condition in Algorithm 2. Figure 1 shows the learning curves of different models on the test data, section 23 of Penn Treebank. Ordinal regression with uneven margins shows great advantage over the same algorithm using even margins. Its performance is also better than perceptron that is trained to separate the best parse from the rest.

By estimating the number of rounds of iterations on the development data, we get the results for the test data as shown in Table 1. Our ordinal regression algorithm with uneven margins achieves the best result in f-score. It verifies that using pairs in training is helpful for the reranking problem. In addition, uneven margins are crucial to using ordinal regression to the reranking task.

In Algorithm 2, we update the weight vector on the sentence level so as to speed up the training, while in Algorithm 1 we update the weight vector for each pair of parses. Figure 3.a shows the comparison of the learning curves of the ordinal regression using parse level updating and the ordinal regression using sentence level updating. Algorithm 2 converges about 40% faster. The performance of Algorithm 2 is very good even within the first few rounds of iterations. Furthermore, the f-score on the test data remains at a high level although it is over-trained. Algorithm 1 easily leads to overfitting for the training data, while Algorithm 2 does not suffer from overfitting. This can be explained by an analog to the gradient methods. For Algorithm 1, we move in one direction at a time, so the result depends on the order of parses of a sentences, and it is easy to jump into a sub-optimum. For Algorithm 2, we move in multiple-directions at a time, so the result is more stable.

Our last set of experiments are about using all and partial pairwise samples. In order to theoretically justify Algorithm 2, we only use $k - 1$ pairwise parses for each sentence, e.g. pairs of parses with consecutive ranks. In Figure 3.b, we compare the results of using all pairs with the results when we use pairs of parses with consecutive ranks. Using only partial pairs makes the algorithm converge much slower.

6 Conclusions

In this paper, we have proposed a general framework for reranking. In this framework, we have proposed two new variants of perceptron. Compared to the previous percep-

section 23, ≤ 100 words (2416 sentences)			
model	recall%	prec%	f-score%
baseline	88.1	88.3	88.2
best-rest	89.2	89.8	89.5
ordinal	88.1	87.8	88.0
uneven ordinal	89.5	90.0	89.8

Table 1. Experimental Results

tron reranking algorithms, the new algorithms use full pairwise samples and allow us to search for margins in a larger space, which are unavailable in the previous works on reranking. We also keep the data complexity unchanged and make the training efficient for these algorithms. Using the new perceptron like algorithms, we investigated the margin selection problem for the parse reranking task. By using uneven margin on ordinal regression, we achieves an f-score of 89.8% on sentences with ≤ 100 words in section 23 of Penn Treebank. The results on margin selection can be employed in reranking systems based on other machine learning algorithms, such as Winnow, Boosting and SVMs.

References

1. Collins, M.: Discriminative reranking for natural language parsing. In: ICML 2000. (2000)
2. Rosenblatt, F.: The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review* **65** (1958) 386–408
3. Novikoff, A.B.J.: On convergence proofs on perceptrons. In: The Symposium on the Mathematical Theory of Automata. Volume 12. (1962)
4. Collins, M., Duffy, N.: New ranking algorithms for parsing and tagging: Kernels over discrete structures, and the voted perceptron. In: Proceedings of ACL 2002. (2002)
5. Ratnaparkhi, A.: A linear observed time statistical parser based on maximum entropy models. In: the Second Conference on Empirical Methods in Natural Language Processing. (1997)
6. Charniak, E.: A maximum-entropy-inspired parser. In: Proceedings of NAACL 2000. (2000)
7. Schapire, R.E., Freund, Y., Bartlett, P., Lee, W.S.: Boosting the margin: a new explanation for the effectiveness of voting methods. In: Proc. 14th International Conference on Machine Learning. (1997) 322–330
8. Shen, L., Joshi, A.K.: An SVM based voting algorithm with application to parse reranking. In: Proc. of CoNLL 2003. (2003)
9. Crammer, K., Singer, Y.: PRanking with Ranking. In: NIPS 2001. (2001)
10. Herbrich, R., Graepel, T., Obermayer, K.: Large margin rank boundaries for ordinal regression. In: Advances in Large Margin Classifiers. MIT Press (2000) 115–132
11. Vapnik, V.N.: Statistical Learning Theory. John Wiley and Sons, Inc. (1998)
12. Zhang, T.: Large Margin Winnow Methods for Text Categorization. In: KDD-2000 Workshop on Text Mining. (2000)
13. Krauth, W., Mezard, M.: Learning algorithms with optimal stability in neural networks. *Journal of Physics A* **20** (1987) 745–752
14. Marcus, M.P., Santorini, B., Marcinkiewicz, M.A.: Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics* **19** (1994) 313–330

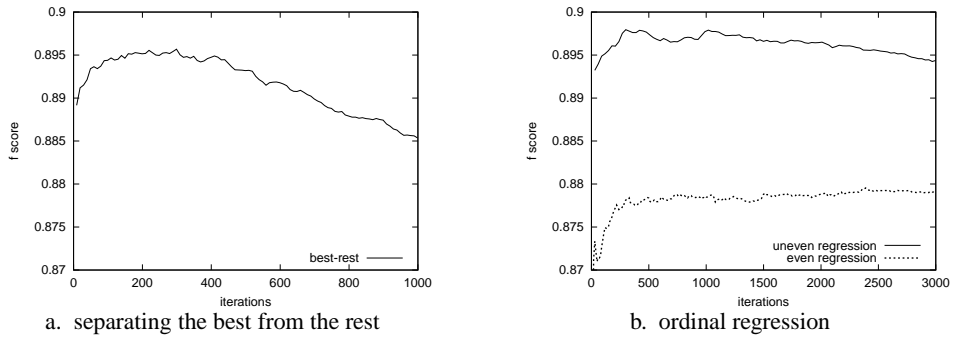


Fig. 1. Learning curves on PTB section 23

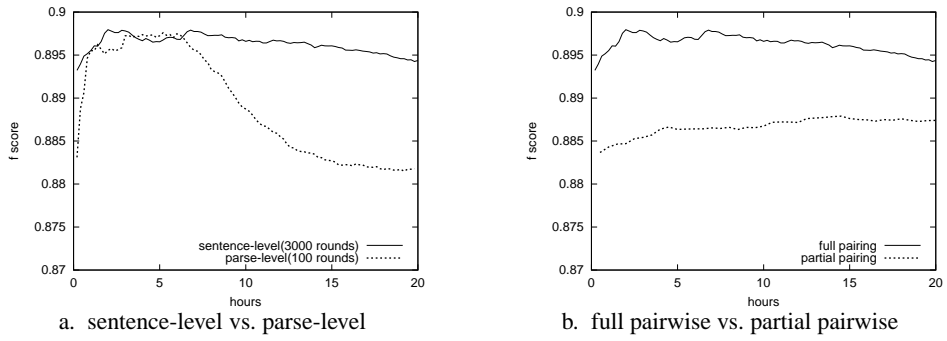


Fig. 2. Learning curves on PTB section 23

15. Freund, Y., Schapire, R.E.: Large margin classification using the perceptron algorithm. *Machine Learning* **37** (1999) 277–296