

# Une méthode discriminant formation simple pour la traduction automatique avec Grands Caractéristiques

Tian Xia Shaodan Zhai Zhongliang Li Shaojun Wang  
Wright State University  
3640 Colonel Glenn Hwy, Dayton, OH 45435  
{Xia.7, Zhai.6, Li.141, Shaojun.Wang}@wright.edu

**Résumé.** Marge infusé algorithmes détendus (MIRAS) dominant modèle de tuning dans la traduction automatique statistique dans le cas des grandes caractéristiques de l'échelle, mais ils sont également célèbres pour la complexité de mise en œuvre. Nous introduisons une nouvelle méthode, qui concerne une liste des N meilleures comme une permutation et minimise la perte Plackett-Luce de permutations rez-de-vérité. Des expériences avec des caractéristiques à grande échelle démontrent que, la nouvelle méthode est plus robuste que MERT ; si ce est seulement à rattacher avec Miras, il a un avantage comparativement, plus facile à mettre en œuvre.

## Abstract.

### A Simple Discriminative Training Method for Machine Translation with Large-Scale Features

The margin infused relaxed algorithm (MIRAS) dominates model tuning in statistical machine translation in the case of large scale features, but also they are famous for the complexity in implementation. We introduce a new method, which regards an N-best list as a permutation and minimizes the Plackett-Luce loss of ground-truth permutations. Experiments with large-scale features demonstrate that, the new method is more robust than MERT ; though it is only matchable with MIRAS, it has a comparatively advantage, easier to implement.

**Mots-clés :** Traduction automatique ; ajustement du modèle ; caractéristiques à grande échelle.

**Keywords:** machine translation ; model tuning ; large-scale features .

## 1 Introduction

Since Och (Och, 2003) proposed minimum error rate training (MERT) to exactly optimize objective evaluation measures, MERT has become a standard model tuning technique in statistical machine translation (SMT). Though MERT performs better by improving its searching algorithm (Macherey *et al.*, 2008; Cer *et al.*, 2008; Galley & Quirk, 2011; Moore & Quirk, 2008; Kumar *et al.*, 2009), it does not work reasonably when there are lots of features<sup>1</sup>. As a result, the margin infused relaxed algorithm (MIRA) dominates in this case (McDonald *et al.*, 2005; Watanabe *et al.*, 2007; Chiang *et al.*, 2008; Tan *et al.*, 2013; Cherry & Foster, 2012).

In SMT, MIRAS consider margin losses related to sentence-level BLEUs. However, since the BLEU is not decomposable into each sentence, these MIRA algorithms use some heuristics to compute the exact losses, e.g., pseudo-document (Chiang *et al.*, 2008), and document-level loss (Tan *et al.*, 2013).

Recently, another successful work in large-scale feature tuning include force decoding based (Yu *et al.*, 2013), classification based (Hopkins & May, 2011).

We aim to provide a simpler tuning method for large-scale features than MIRAS. Our motivation derives from an observation on MERT. As MERT considers the quality of only top1 hypothesis set, there might have more-than-one set of parameters, which have similar top1 performances in tuning, but have very different topN hypotheses. Empirically, we expect an ideal model to benefit the total N-best list. That is, better hypotheses should be assigned with higher ranks, and this might decrease the error risk of top1 result on unseen data.

Plackett (Plackett, 1975) offered an easy-to-understand theory of modeling a permutation. An N-best list is assumedly generated by sampling without replacement. The *i*th hypothesis to sample relies on those ranked after it, instead of on the

1. The regularized MERT seems promising from Galley et al. (Galley *et al.*, 2013) at the cost of model complexity.

whole list. This model also supports a partial permutation which accounts for top  $k$  positions in a list, regardless of the remaining. When taking  $k$  as 1, this model reduces to a standard conditional probabilistic training, whose dual problem is actual the maximum entropy approach (Och & Ney, 2002; Berger *et al.*, 1996). Although Och (Och, 2003) substituted direct error optimization for a maximum entropy based training, probabilistic models correlate with BLEU well when features are rich enough. The similar claim also appears in (Zhu & Hastie, 2001). This also make the new method be applicable in large-scale features.

## 2 Plackett-Luce Model

Plackett-Luce was firstly proposed to predict ranks of horses in gambling (Plackett, 1975). Let  $\mathbf{r} = (r_1, r_2 \dots r_N)$  be  $N$  horses with a probability distribution  $\mathcal{P}$  on their abilities to win a game, and a rank  $\pi = (\pi(1), \pi(2) \dots \pi(|\pi|))$  of horses can be understood as a generative procedure, where  $\pi(j)$  denotes the index of the horse in the  $j$ th position.

In the 1st position, there are  $N$  horses as candidates, each of which  $r_j$  has a probability  $p(r_j)$  to be selected. Regarding the rank  $\pi$ , the probability of generating the champion is  $p(r_{\pi(1)})$ . Then the horse  $r_{\pi(1)}$  is removed from the candidate pool.

In the 2nd position, there are only  $N - 1$  horses, and their probabilities to be selected become  $p(r_j)/Z_2$ , where  $Z_2 = 1 - p(r_{\pi(1)})$  is the normalization. Then the runner-up in the rank  $\pi$ , the  $\pi(2)$ th horse, is chosen at the probability  $p(r_{\pi(2)})/Z_2$ . We use a consistent terminology  $Z_1$  in selecting the champion, though  $Z_1$  equals 1 trivially.

This procedure iterates to the last rank in  $\pi$ . The key idea for the Plackett-Luce model is the choice in the  $i$ th position in a rank  $\pi$  only depends on the candidates not chosen at previous stages. The probability of generating a rank  $\pi$  is given as follows

$$p(\pi) = \prod_{j=1}^{|\pi|} \frac{p(r_{\pi(j)})}{Z_j} \quad (1)$$

where  $Z_j = 1 - \sum_{t=1}^{j-1} p(r_{\pi(t)})$ .

We offer a toy example (Table 1) to demonstrate this procedure.

$\mathbf{r}$	$r_1$	$r_2$	$r_3$
$\pi$	2	3	1
$Z$	1	$1-p(r_2)$	$1-(p(r_2) + p(r_3))$
$p(\pi)$	$\frac{p(r_2)}{Z_1}$	$\frac{p(r_3)}{Z_2}$	$\frac{p(r_1)}{Z_3}$

TABLE 1 – The probability of the rank  $\pi = (2, 3, 1)$  is  $p(r_2) \cdot p(r_3)/(1 - p(r_2))$  in a simplified form, as  $\frac{p(r_2)}{Z_1} = p(r_2)$  and  $\frac{p(r_1)}{Z_3} = 1$ .

**Theorem 1** *The permutation probabilities  $p(\pi)$  form a probability distribution over a set of permutations  $\Omega_\pi$ . For example, for each  $\pi \in \Omega_\pi$ , we have  $p(\pi) > 0$ , and  $\sum_{\pi \in \Omega_\pi} p(\pi) = 1$ .*

We have to note that,  $\Omega_\pi$  is not necessarily required to be completely ranked permutations in theory and in practice, since gamblers might be interested in only the champion and runner-up, and thus  $|\pi| \leq N$ . In experiments, we would examine the effects on different length of permutations, systems being termed  $PL(|\pi|)$ .

**Theorem 2** *Given any two permutations  $\pi$  and  $\pi'$ , and they are different only in two positions  $p$  and  $q$ ,  $p < q$ , with  $\pi(p) = \pi'(q)$  and  $\pi(q) = \pi'(p)$ . If  $p(\pi(p)) > p(\pi(q))$ , then  $p(\pi) > p(\pi')$ .*

In other words, exchanging two positions in a permutation where the horse more likely to win is not ranked before the other would lead to an increase of the permutation probability.

This suggests the ground-truth permutation, ranked decreasingly by their probabilities, owns the maximum permutation probability on a given distribution. In SMT, we are motivated to optimize parameters to maximize the likelihood of ground-truth permutation of an N-best hypotheses.

Due to the limitation of space, see (Plackett, 1975; Cao *et al.*, 2007) for the proofs of the theorems.

### 3 Plackett-Luce Model in Statistical Machine Translation

In SMT, let  $\mathbf{f} = (f_1, f_2 \dots)$  denote source sentences, and  $\mathbf{e} = (\{e_{1,1}, \dots\}, \{e_{2,1}, \dots\} \dots)$  denote target hypotheses. A set of features are defined on both source and target side. We refer to  $h(e_{i,*})$  as a feature vector of a hypothesis from the  $i$ th source sentence, and its score from a ranking function is defined as the inner product  $h(e_{i,*})^T \mathbf{w}$  of the weight vector  $\mathbf{w}$  and the feature vector.

We first follow the popular exponential style to define a parameterized probability distribution over a list of hypotheses.

$$p(e_{i,j}) = \frac{\exp\{h(e_{i,j})^T \mathbf{w}\}}{\sum_k \exp\{h(e_{i,k})^T \mathbf{w}\}} \quad (2)$$

The ground-truth permutation of an  $n$ best list is simply obtained after ranking by their sentence-level BLEUs. Here we only concentrate on their relative ranks which are straightforward to compute in practice, e.g. add 1 smoothing. Let  $\pi_i^*$  be the ground-truth permutation of hypotheses from the  $i$ th source sentences, and our optimization objective is maximizing the log-likelihood of the ground-truth permutations and penalized using a zero-mean and unit-variance Gaussian prior. This results in the following objective and gradient :

$$\mathcal{L} = \log\left\{\prod_i p(\pi_i^*, \mathbf{w})\right\} - \frac{1}{2} \mathbf{w}^T \mathbf{w} \quad (3)$$

$$\frac{\partial \mathcal{L}}{\partial \mathbf{w}} = \sum_i \sum_j \{h(e_{i,\pi_i^*(j)}) - \sum_{t=j} (h(e_{i,\pi_i^*(t)}) \cdot \frac{p(e_{i,\pi_i^*(t)})}{Z_{i,j}})\} - \mathbf{w} \quad (4)$$

where  $Z_{i,j}$  is defined as the  $Z_j$  in Formula (1) of the  $i$ th source sentence.

The log-likelihood function is smooth, differentiable, and concave with the weight vector  $\mathbf{w}$ , and its local maximal solution is also a global maximum. Iteratively selecting one parameter in  $\alpha$  for tuning in a line search style (or MERT style) could also converge into the global maximum (Bertsekas, 1999). In practice, we use the faster limited-memory BFGS (L-BFGS) algorithm (Byrd *et al.*, 1995).

#### N-best Hypotheses Resample

The log-likelihood of a Plackett-Luce model is not a strict upper bound of the BLEU score, however, it correlates with BLEU well in the case of rich features. The concept of ‘‘rich’’ is actually qualitative, and obscure to define in different applications. We empirically provide a formula to measure the richness in the scenario of machine translation.

$$r = \frac{\text{the size of features}}{\text{the average size of N-best lists}} \quad (5)$$

The greater, the richer. In practice, we find a rough threshold of  $r$  is 5.

In engineering, the size of an N-best list with unique hypotheses is usually less than several thousands. This suggests that, if features are up to thousands or more, the Plackett-Luce model is quite suitable here. Otherwise, we could reduce the size of N-best lists by sampling to make  $r$  beyond the threshold.

There may be other efficient sampling methods, and here we adopt a simple one. If we want to  $m$  samples from a list of hypotheses  $\mathbf{e}$ , first, the  $\frac{m}{3}$  best hypotheses and the  $\frac{m}{3}$  worst hypotheses are taken by their sentence-level BLEUs. Second, we sample the remaining hypotheses on distribution  $p(e_i) \propto \exp(h(e_i)^T \mathbf{w})$ , where  $\mathbf{w}$  is an initial weight from last iteration.

## 4 Evaluation

	MT02(dev)	MT04(%)	MT05(%)
MERT	34.61	31.76	28.85
MIRA	35.31	32.25	29.37
PL(1)	34.20	31.70	28.90
PL(2)	34.31	31.83	29.10
PL(3)	34.39	32.05	29.20
PL(4)	34.40	32.13	29.46+
PL(5)	34.46	32.19+	29.42+
PL(6)	34.37	32.16	29.30
PL(7)	34.39	32.20+	29.32
PL(8)	34.70	32.19+	29.10
PL(9)	34.30	32.07	29.22
PL(10)	34.30	32.14	29.19

TABLE 2 – PL( $k$ ) : Plackett-Luce model optimizing the ground-truth permutation with length  $k$ . The significant symbols (+ at 0.05 level) are compared with MERT.

We compare our method with MERT and MIRA<sup>2</sup> in two tasks, iterative training, and N-best list reranking. We do not list PRO (Hopkins & May, 2011) as our baseline, as Cherry et al. (Cherry & Foster, 2012) have compared PRO with MIRA and MERT massively.

In the first task, we align the FBIS data (about 230K sentence pairs) with GIZA++, and train a 4-gram language model on the Xinhua portion of Gigaword corpus. A hierarchical phrase-based (HPB) model (Chiang, 2007) is tuned on NIST MT 2002, and tested on MT 2004 and 2005. All features are eight basic ones (Chiang, 2007) and extra 220 group features. We design such feature templates to group grammar rules by the length of source side and target side, (feat-type,  $a \leq \text{src-side} \leq b, c \leq \text{tgt-side} \leq d$ ), where the feat-type denotes any of the relative frequency, reversed relative frequency, lexical probability and reversed lexical probability, and [a, b], [c, d] enumerate all possible subranges of [1, 10], as the maximum length on both sides of a hierarchical grammar is limited to 10. There are  $4 \times 55$  extra group features.

In the second task, we rerank an N-best list from a HPB system with 7491 features from a third party. The system uses six million parallel sentence pairs available to the DARPA BOLT Chinese-English task. This system includes 51 dense features (translation probabilities, provenance features, etc.) and up to 7440 sparse features (mostly lexical and fertility-based). The language model is a 6-gram model trained on a 10 billion words, including the English side of our parallel corpora plus other corpora such as Gigaword (LDC2011T07) and Google News. For the tuning and test sets, we use 1275 and 1239 sentences respectively from the LDC2010E30 corpus.

### 4.1 Plackett-Luce Model for SMT Tuning

We conduct a full training of machine translation models. By default, a decoder is invoked for at most 40 times, and each time it outputs 200 hypotheses to be combined with those from previous iterations and sent into tuning algorithms.

In getting the ground-truth permutations, there are many ties with the same sentence-level BLEU, and we just take one randomly. In this section, all systems have only around two hundred features, hence in Plackett-Luce based training, we sample 30 hypotheses in an accumulative  $n$ best list in each round of training.

All results are shown in Table 2, we can see that all PL( $k$ ) systems do not perform well as MERT or MIRA in the development data, this maybe due to that PL( $k$ ) systems do not optimize BLEU and the features here are relatively not enough compared to the size of N-best lists (empirical Formula 5). However, PL( $k$ ) systems are better than MERT in testing. PL( $k$ ) systems consider the quality of hypotheses from the 2th to the  $k$ th, which is guessed to act the role of the margin like SVM in classification. Interestingly, MIRA wins first in training, and still performs quite well in testing.

The PL(1) system is equivalent to a max-entropy based algorithm (Och & Ney, 2002) whose dual problem is actually maximizing the conditional probability of the best BLEU hypothesis. When we increase the  $k$ , performances improve at first. After reaching a maximum around  $k = 5$ , they decrease slowly. We explain this phenomenon as this, when features

2. MIRA is from the open-source Moses (Koehn et al., 2007)

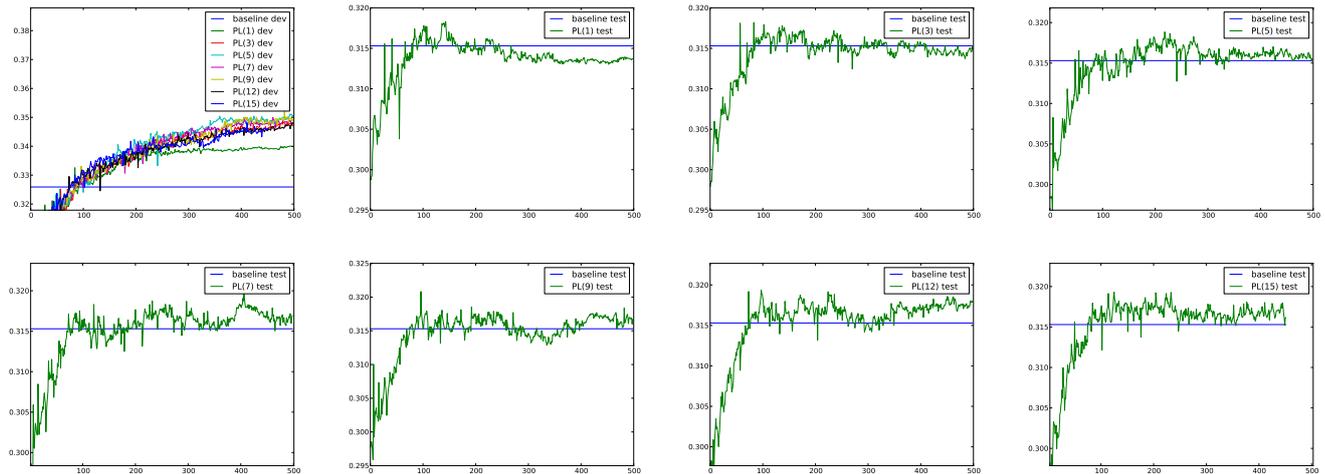


FIGURE 1 –  $PL(k)$  with 500 L-BFGS iterations,  $k=1,3,5,7,9,12,15$  compared with MIRA in reranking.

are rich enough, higher BLEU scores could be easily fitted, then longer ground-truth permutations include more useful information.

## 4.2 Plackett-Luce Model for SMT Reranking

After being de-duplicated, the N-best list has an average size of around 300, and with 7491 features. Referring to Formula 5, this is ideal to use the Plackett-Luce model. Results are shown in Figure 1. We observe some interesting phenomena.

First, the Plackett-Luce models boost the training BLEU very greatly, even up to 2.5 points higher than MIRA. This verifies our assumption, richer features benefit BLEU, though they are optimized towards a different objective.

Second, the over-fitting problem of the Plackett-Luce models  $PL(k)$  is alleviated with moderately large  $k$ . In  $PL(1)$ , the over-fitting is quite obvious, the portion in which the curve overpasses MIRA is the smallest compared to other  $k$ , and its convergent performance is below the baseline. When  $k$  is not smaller than 5, the curves are almost above the MIRA line. After 500 L-BFGS iterations, their performances are no less than the baseline, though only by a small margin.

This experiment displays, in large-scale features, the Plackett-Luce model correlates with BLEU score very well, and alleviates overfitting in some degree.

## 5 Conclusion

This work has successfully brought the Plackett-Luce loss into statistical machine translation for model tuning. When in the case of large scale features, our method works better than MERT, and similarly with MIRA. This method, compared with MIRA, has a merit in practice, that is simpler to implement. Moreover, in the N-best list reranking experiment, our method also shows interesting property, resilient to overfitting. In the future work, we plan to expand the Plackett-Luce loss on packed translation forests to acquire better performance.

## Références

- BERGER A. L., PIETRA V. J. D. & PIETRA S. A. D. (1996). A maximum entropy approach to natural language processing. *Computational linguistics*, **22**(1).
- BERTSEKAS D. P. (1999). Nonlinear programming.
- BYRD R. H., LU P., NOCEDAL J. & ZHU C. (1995). A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, **16**(5).

- CAO Z., QIN T., LIU T.-Y., TSAI M.-F. & LI H. (2007). Learning to rank : from pairwise approach to listwise approach. In *Proc. of ICML*.
- CER D., JURAFSKY D. & MANNING C. D. (2008). Regularization and search for minimum error rate training. In *Proc. of WSMT*.
- CERRY C. & FOSTER G. (2012). Batch tuning strategies for statistical machine translation. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics : Human Language Technologies*, p. 427–436 : Association for Computational Linguistics.
- CHIANG D. (2007). Hierarchical phrase-based translation. *computational linguistics*, **33**(2).
- CHIANG D., MARTON Y. & RESNIK P. (2008). Online large-margin training of syntactic and structural translation features. In *Proc. of EMNLP*.
- GALLEY M. & QUIRK C. (2011). Optimal search for minimum error rate training. In *Proc. of EMNLP*.
- GALLEY M., QUIRK C., CERRY C. & TOUTANOVA K. (2013). Regularized minimum error rate training. In *EMNLP*, p. 1948–1959.
- HOPKINS M. & MAY J. (2011). Tuning as ranking. In *Proc. of EMNLP*.
- KOEHN P., HOANG H., BIRCH A., CALLISON-BURCH C., FEDERICO M., BERTOLDI N., COWAN B., SHEN W., MORAN C., ZENS R. *et al.* (2007). Moses : Open source toolkit for statistical machine translation. In *Proc. of ACL : Poster*.
- KUMAR S., MACHEREY W., DYER C. & OCH F. (2009). Efficient minimum error rate training and minimum bayes-risk decoding for translation hypergraphs and lattices. In *Proc. of Joint ACL and AFNLP*.
- MACHEREY W., OCH F. J., THAYER I. & USZKOREIT J. (2008). Lattice-based minimum error rate training for statistical machine translation. In *Proc. of EMNLP*.
- MCDONALD R., CRAMMER K. & PEREIRA F. (2005). Online large-margin training of dependency parsers. In *Proc. of ACL*.
- MOORE R. C. & QUIRK C. (2008). Random restarts in minimum error rate training for statistical machine translation. In *Proceedings of the 22nd International Conference on Computational Linguistics-Volume 1*, p. 585–592 : Association for Computational Linguistics.
- OCH F. J. (2003). Minimum error rate training in statistical machine translation. In *Proc. of ACL*.
- OCH F. J. & NEY H. (2002). Discriminative training and maximum entropy models for statistical machine translation. In *Proc. of ACL*.
- PLACKETT R. L. (1975). The analysis of permutations. *Applied Statistics*.
- TAN M., XIA T., WANG S. & ZHOU B. (2013). A corpus level mira tuning strategy for machine translation. In *EMNLP*, p. 851–856.
- WATANABE T., SUZUKI J., TSUKADA H. & ISOZAKI H. (2007). Online large-margin training for statistical machine translation. In *Proc. of EMNLP*.
- YU H., HUANG L., MI H. & ZHAO K. (2013). Max-violation perceptron and forced decoding for scalable mt training. In *EMNLP*, p. 1112–1123.
- ZHU J. & HASTIE T. (2001). Kernel logistic regression and the import vector machine. In *Proc. of NIPS*.