

An Improved Chinese Word Segmentation System with Conditional Random Field

Hai Zhao, Chang-Ning Huang and Mu Li

Microsoft Research Asia,

49, Zhichun Road, Haidian District, Beijing, P. R. China, 100080

Email: {f-hzhao,cnhuang}@msrchina.research.microsoft.com,
muli@microsoft.com

Abstract

In this paper, we describe a Chinese word segmentation system that we developed for the Third SIGHAN Chinese Language Processing Bakeoff (Bakeoff-2006). We took part in six tracks, namely the closed and open track on three corpora, Academia Sinica (CKIP), City University of Hong Kong (CityU), and University of Pennsylvania/University of Colorado (UPUC). Based on a conditional random field based approach, our word segmenter achieved the highest F measures in four tracks, and the third highest in the other two tracks. We found that the use of a 6-tag set, tone feature of Chinese character and assistant segmenters trained on other corpora further improve Chinese word segmentation performance.

1 Introduction

Conditional random field (CRF) is a statistical sequence modeling framework first introduced into language processing in (Lafferty et al., 2001). In (Peng et al., 2004), this framework is used for Chinese word segmentation by treating it as a binary decision task, such that each Chinese character is labeled either as the beginning of a word or not.

Since two participants, (Low et al., 2005) and (Tseng et al., 2005) in Bakeoff-2005, have given the best results in almost all word segmentation tracks, we continue to improve CRF-based tagging method

for Chinese word segmentation on their track. Our implementation used CRF++ package Version 0.41¹ by Taku Kudo.

In our system, a character in the given sequence is labeled by a tag which stands for its position in the word that the character belongs to. We handle closed and open test in the same way. The difference is that those features concerned with additional linguistic resources are added in the feature set of closed test to produce the feature set for open test.

2 Tag Set Selection

Character based tagging method for Chinese word segmentation, either based on maximum entropy or CRF, views Chinese word segmentation as a label tagging problem, which is described in detail in (Ratnaparkhi, 1996).

The probability model and corresponding feature function is defined over the set $H \times T$, where H is the set of possible contexts (or any predefined condition) and T is the set of possible tags. Generally, a feature function can be defined as follows,

$$f(h, t) = \begin{cases} 1, & \text{if } h = h_i \text{ is satisfied and } t = t_j \\ 0, & \text{otherwise,} \end{cases} \quad (1)$$

where $h_i \in H$ and $t_j \in T$.

For convenience, features are generally organized into some groups, which used to be called feature templates. For example, a bigram feature template C_1 stands for the next character occurring in the corpus after each character.

¹<http://chasen.org/~taku/software/CRF++/>

Table 1: Feature templates

Code	Type	Feature	Function
a	Unigram	C_{-1}, C_0, C_1	The previous, current, and next character
b	Bigram	$C_{-1}C_0, C_0C_1$	The previous (next) and current characters
c	Jump	$C_{-1}C_1$	The previous and next characters
d	Punctuation	$Pu(C_0)$	Current character is a punctuation or not
e	Date, Digital and Letter	$T_{-1}T_0T_1$	Types of previous, current and next character
f	Tone	$To(C_0)$	Tone of current character

As for tag set, there are two kinds of schemes that are used to distinguish the character position in a word in the previous work, i.e., 4-tag set and 2-tag set. The details are listed in Table 2. Note that a 4-tag set is used for maximum entropy model in (Xue, 2003; Xue and Shen, 2003) and (Low et al., 2005), while a 2-tag set is used for CRF model in (Peng et al., 2004) and (Tseng et al., 2005).

Table 2: Tag sets in the previous work

4-tag set Xue/(Low)		2-tag set Peng/Tseng	
Tag	Function	Tag	Function
B(LL)	begin	Start	start
M(MM)	middle	NoStart	continuation
E(RR)	end		
S(LR)	single		

Generally speaking, activated feature functions in practice like (1) are determined by both feature template and tag set. In the existing work, tag set is specified beforehand. To effectively perform tagging for those long words, we extend the 4-tag set of Xue/Low into a 6-tag set. Two tags, ‘ B_2 ’ and ‘ B_3 ’, are added into a 4-tag set to form a 6-tag set, which stands for the second and the third character position in a word, respectively.

3 Feature Templates for Closed Test

The feature template set we selected for closed test is shown in Table 1. Note that only six n -gram feature templates are used in our system rather than more than ten ones in previous work. Here we give an explanation to feature template (e) and (f).

Feature template (e) is improved from the corresponding one in (Low et al., 2005). T_n , where $n =$

$-1, 0, 1$, stands for four predefined class: numbers represent class 1, those characters, whose meanings are dates and time such as ‘年’, ‘月’, and ‘秒’ etc, represent class 2, English letters represent class 3, and other characters represent class 4.

As for feature template (f), $To(C_0)$ stands for the tone of current character. There are five possible types of tones for Chinese characters in mandarin, we just assign 0, 1, 2, 3 and 4 as feature values. For example, consider some characters, ‘中’, ‘国’, ‘很’, ‘大’ and ‘吗’, $To(C_0)$ is 1, 2, 3, 4 and 0, respectively.

4 Feature Templates for Open Test

In open test, we use two kinds of extra feature templates to improve the performance upon closed test.

4.1 External Dictionary

External dictionary features are introduced in (Low et al., 2005). We continue to use the online dictionary from Peking University downloadable from the Internet², consisting of about 108,000 words of one to four characters. If there is some subsequences of neighboring characters around C_0 in the sentence that match words in this dictionary, then the longest one W in the dictionary will be chosen. The following feature templates will be added:

(g) Lt_0

(h) $C_n t_0 (n = -1, 0, 1)$

where t_0 is the boundary tag of C_0 in W , and L is the number of characters in W .

²[http://ccl.pku.edu.cn/doubtfire/Course/Chinese %20Information%20Processing/Source Code/ Chapter 8/Lexicon full 2000.zip](http://ccl.pku.edu.cn/doubtfire/Course/Chinese%20Information%20Processing/Source%20Code/Chapter%208/Lexicon%20full%202000.zip)

4.2 Assistant Segmenter

We observed that although different segmentation standards exist, they share the same way for most word segmentations. Thus, though those segmenters trained on different corpora will give some different segmentation results, they agree on most cases. In fact, we find that it is feasible to customize a pre-defined standard into any other standards with TBL method in (Gao et al., 2005). And it is also worth incorporating different segmenters into one segmenter based on the current standard. For convenience, we call the segmenter subjected to the current standard main segmenter, and the others assistant segmenters.

A feature template will be added for an assistant segmenter:

- (i) $t(C_0)$

where $t(C_0)$ is the output tag of the assistant segmenter for the current character C_0 . For example, consider character sequence, '我们都是中国人', an assistant segmenter gives the tag sequence 'BESSBES' according to its output segmentation, then $t(C_0)$ by this assistant segmenter is 'B', 'E', 'S', 'S', 'B', 'E', and 'S' for each current character, respectively.

In our system, we integrate all other segmenters that are trained on all corpora from Bakeoff-2003, 2005 and 2006 with the feature set used in closed test (Sproat and Emerson, 2003; Emerson, 2005). The segmenter, MSRSeg, described in (Gao et al., 2003) is also integrated.

Our assistant segmenter method is more convenient compared to the additional training corpus method in (Low et al., 2005). Firstly, the performance of additional corpus method depends on the performance of the trained segmenter that carries out the corpus extraction task. If the segmenter is not well-trained, then it cannot effectively extract the most wanted additional corpus to some extent. Secondly, additional corpus method is only able to integrate useful corpus, but it cannot integrate a well-trained segmenter while the corpus cannot be accessed. Finally, additional corpus method is very difficult to use in CRF model, the reason is that the increase of corpus can lead to a dramatic increase of memory and time consuming in this case, while assistant segmenters just lead to little increase of memory and time consuming in training.

It is more interesting that we may also regard the external dictionary method as another assistant segmenter in some degree, that is, a maximal matching segmenter with the specified external dictionary. Thus, all of our additional methods in open test can be viewed as assistant segmenter ones.

5 Evaluation Results

We took part in six segmentation tasks in Bakeoff-2006, namely the closed and open track on three corpora, Academia Sinica (CKIP), City University of Hong Kong (CityU), and University of Pennsylvania/University of Colorado (UPUC). The default metric of segmentation performance is F measure.

The comparisons of official results between ours and the other best of Bakeoff-2006 are in Table 3.

Table 3: Comparisons of the other best results of Bakeoff-2006 and ours

Type	Participant	Corpora		
		CKIP	CityU	UPUC
Closed Test	Bakeoff-2006	0.957	0.972	0.930
	Our results	0.958	0.971	0.933
Open Test	Bakeoff-2006	0.954	0.976	0.944
	Our results	0.959	0.977	0.933

Our system achieved the highest F measures in four tracks, and the third highest in the other two tracks. However, a format error unfortunately occurred in the open test of UPUC corpus as we submitted our final results. Thus an abnormal result in this task is obtained, the official F measure in open test is the same as that in closed test. We get the actual F measure of 0.953 after the bug is fixed.

The results in MSRA corpus from our evaluation are listed in Table 4.

Table 4: Comparison of the best official results of Bakeoff-2006 and ours on MSRA corpora

Type	Bakeoff-2006	Ours
Closed Test	0.963	0.961
Open Test	0.979	0.982

The sizes of training corpora (in number of characters) and difference of our results between open

and closed test are shown in Table 5. This illustrates how much assistant segmenters improve the segmentation performance for different training corpora, and how the size of training corpus affects the improvement contributed by assistant segmenters.

Table 5: The sizes of training corpora and difference of our results between open test and closed test

	CKIP	CityU	MSRA	UPUC
Difference	0.001	0.006	0.021	0.020
Size	9M	2.9M	2.3M	0.88M

References

- Thomas Emerson. 2005. The second international Chinese word segmentation bakeoff. In *Proceedings of the Fourth SIGHAN Workshop on Chinese Language Processing*, pages 123–133, Jeju Island, Korea.
- Jianfeng Gao, Mu Li, and Chang-Ning Huang. 2003. Improved source-channel models for Chinese word segmentation. In *Proceedings 41st Annual Meeting of the Association for Computational Linguistics*, pages 272–279, Sapporo, Japan, July 7-12.
- Jianfeng Gao, Mu Li, Andi Wu, and Chang-Ning Huang. 2005. Chinese word segmentation and named entity recognition: A pragmatic approach. *Computational Linguistics*, 31(4):531–574.
- John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *ICML '01: Proceedings of the Eighteenth International Conference on Machine Learning*, pages 282–289, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Jin Kiat Low, Hwee Tou Ng, and Wenyuan Guo. 2005. A maximum entropy approach to Chinese word segmentation. In *Proceedings of the Fourth SIGHAN Workshop on Chinese Language Processing*, pages 161–164, Jeju Island, Korea.
- Fuchun Peng, Fangfang Feng, and Andrew McCallum. 2004. Chinese segmentation and new word detection using conditional random fields. In *COLING 2004*, pages 562–568, Geneva, Switzerland, August 23-27.
- Adwait Ratnaparkhi. 1996. A maximum entropy part-of-speech tagger. In *Proceedings of the Empirical Method in Natural Language Processing Conference*, pages 133–142, University of Pennsylvania.
- Richard Sproat and Thomas Emerson. 2003. The first international Chinese word segmentation bakeoff. In *The Second SIGHAN Workshop on Chinese Language Processing*, pages 133–143, Sapporo, Japan.
- Huihsin Tseng, Pichuan Chang, Galen Andrew, Daniel Jurafsky, and Christopher Manning. 2005. A conditional random field word segmenter for SIGHAN bakeoff 2005. In *Proceedings of the Fourth SIGHAN Workshop on Chinese Language Processing*, pages 168–171, Jeju Island, Korea.
- Nianwen Xue and Libin Shen. 2003. Chinese word segmentation as LMR tagging. In *Proceedings of the Second SIGHAN Workshop on Chinese Language Processing, in conjunction with ACL'03*, pages 176–179, Sapporo, Japan.
- Nianwen Xue. 2003. Chinese word segmentation as character tagging. *Computational Linguistics and Chinese Language Processing*, 8(1):29–48.