

Does Higher Order LSTM Have Better Accuracy for Segmenting and Labeling Sequence Data?

Yi Zhang, Xu Sun, Shuming Ma, Yang Yang, Xuancheng Ren

MOE Key Lab of Computational Linguistics, School of EECS, Peking University
{zhangyi16, xusun, shumingma, 1200012760, renxc}@pku.edu.cn

Abstract

Existing neural models usually predict the tag of the current token independent of the neighboring tags. The popular LSTM-CRF model considers the tag dependencies between every two consecutive tags. However, it is hard for existing neural models to take longer distance dependencies of tags into consideration. The scalability is mainly limited by the complex model structures and the cost of dynamic programming during training. In our work, we first design a new model called “high order LSTM” to predict multiple tags for the current token which contains not only the current tag but also the previous several tags. We call the number of tags in one prediction as “order”. Then we propose a new method called Multi-Order BiLSTM (MO-BiLSTM) which combines low order and high order LSTMs together. MO-BiLSTM keeps the scalability to high order models with a pruning technique. We evaluate MO-BiLSTM on all-phrase chunking and NER datasets. Experiment results show that MO-BiLSTM achieves the state-of-the-art result in chunking and highly competitive results in two NER datasets.¹

1 Introduction

Chunking and named entity recognition are sequence labeling tasks whose target is to find the correct segments and give them the correct labels. The tags inside a segment have internal dependencies. The tags in consecutive segments may have dependencies, too. Therefore, it is natural to take the tag dependencies into consideration when making a prediction in such sequence labeling tasks.

Recently, methods have been proposed to capture tag dependencies for neural networks. Collobert et al. (2011) proposed a method based on convolutional neural networks, which can use dynamic programming in training and testing stage (like a CRF layer) to capture tag dependencies. Furthermore, Huang et al. (2015) proposed LSTM-CRF by combining LSTM and CRF for structured learning. They use a transition matrix to model the tag dependencies. A similar structure is adopted by Ma and Hovy (2016). Their model also involves an external layer to extract some character level features.

However, it is not explicit how to model the dependencies of more tags or use the dependency information in these lines of work. We then propose a solution to capture long distance tag dependencies and use them for dependency-aware prediction of tags. For clarity, we first give some detailed explanations of the related terms in our work. “order” means the number of tags that a prediction involves in a model. An order-2 tag is a bigram which contains the previous tag and the current tag at a certain time step, as shown in Figure 1. Higher order tags are defined in a similar way.

We first develop a simple method to implement high order models. But these models, which are supposed to capture more tag dependency information, perform worse and worse as the order of models increases. One possible reason is that trying to capture more tag dependencies raises the difficulty of prediction. We name these models as single order models and propose a new method based on them. The proposed **Multi-Order LSTM (MO-LSTM)** combines multi-order information from these single order models to decode. It keeps the scalability with a proposed pruning technique and performs well in our

¹The code is available at <https://github.com/lancopku/Multi-Order-LSTM>
This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>

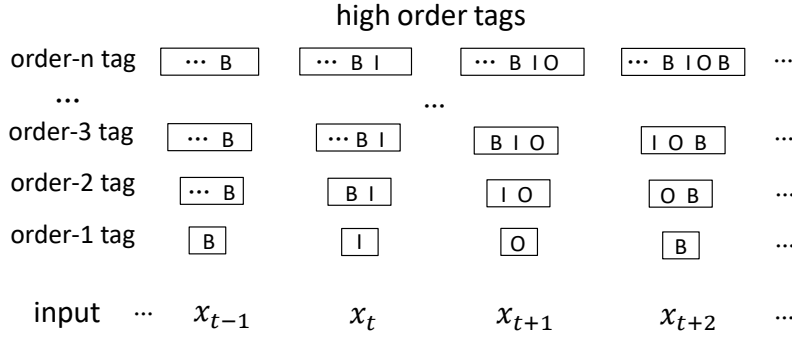


Figure 1: An illustration of tags of different orders.

tasks. Experiments show that MO-LSTM achieves the state-of-the-art F1 score in all-phrase chunking and competitive scores in two NER datasets.

The contributions of this work are as follows:

- We extend the LSTM model to higher order models. However, the performance of the high order models which are supposed to capture longer tag dependencies is getting worse when increasing the order.
- We propose a model integrating low order and high order models. It keeps the scalability in both training and testing stage with a pruning technique.
- The proposed MO-LSTM achieves an evident error reduction in chunking and NER tasks. It produces the state-of-the-art F1 score in chunking and highly competitive results in two NER datasets.

2 Single Order LSTM

We first propose a simple training and decoding method which enables the existing models to extend to higher order models. Take the order-2 model as an example, for each word we combine its previous tag and its current tag to produce a bigram tag as its new tag to predict. Hence, the model can be trained with the “new” bigram (order-2) tag set.

Formally, given an input sequence $\mathbf{x} = \{x_1, x_2, \dots, x_T\}$, where x_t denotes the t -th word in a sentence and T denotes the sentence length. The sequence $\mathbf{y} = \{y_1, y_2, \dots, y_T\}$ represents a possible label sequence for \mathbf{x} . We denote $\mathcal{Y}^{(1)}$ as the set of all possible order-1 labels, and $y_t \in \mathcal{Y}^{(1)}$. The order-1 model can be represented as:

$$s_1(y_1, y_2, \dots, y_T | \mathbf{x}; \theta) = \prod_{t=1}^T s(y_t | \mathbf{x}; \theta) \tag{1}$$

where θ is the parameters of the model. In implementation, we use a Bi-LSTM with a softmax layer to compute the score $s(y_t | \mathbf{x}; \theta)$.

To extend the order-1 model to an order-2 model, we transform the unigram label sequence into a bigram label sequence $y_0 y_1, y_1 y_2, \dots, y_{T-1} y_T$, where y_0 is a special START symbol. The bigram label is defined as a combination of two consecutive label y_{t-1} and y_t , and $\mathcal{Y}^{(2)}$ is the set of all possible bigram labels that appear in the training set. The order-2 model can then be written as:

$$s_2(y_1, y_2, \dots, y_T | \mathbf{x}; \theta) = \prod_{t=1}^T s(y_{t-1} y_t | \mathbf{x}; \theta) \tag{2}$$

Similar to the order-1 model, the score $s(y_{t-1} y_t | \mathbf{x}; \theta)$ is computed by a Bi-LSTM with a softmax layer. In implementation, the difference with the order-1 model is that the unigram label is replaced with the

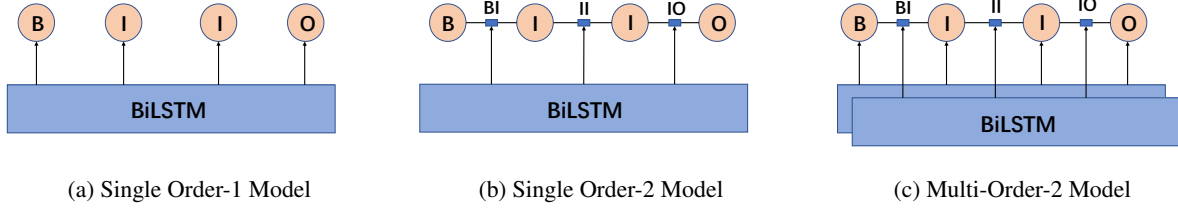


Figure 2: An illustration of the single order model and the multi-order model. The single order-1 model is a BiLSTM.

bigram label. In this way, the model can be further extended to order- n :

$$s_n(y_1, y_2, \dots, y_T | \mathbf{x}; \theta) = \prod_{t=1}^T s(y_{t-n+1} \dots y_t | \mathbf{x}; \theta) \quad (3)$$

As the order of the models increases, the models are supposed to learn more tag dependencies. However, according to our experiments, the performance of these models is getting worse, and the detailed results are shown in Section 4. An intuitive reason to explain the experimental phenomena is that the increasing size of the label set makes it more difficult to predict a correct label of the input word. Another potential reason is that the complex structure leads to overfitting problem. Sun (2014) suggests that complex structures are actually harmful to the generalization ability in structured prediction.

3 Multi-Order BiLSTM

The performance of single high order models deteriorates as the order increases. But they might capture some kinds of useful dependency information. To make use of these dependency information, we introduce a multi-order model which combines the low-order and high-order information. The proposed multi-order model consists of several single order models (as described in Section 2) of different orders. At the training stage, these models are trained separately as usual. At the decoding stage, we propose a new decoding method to combine the low order model and the high order model. Since both low order information and high order information is used when decoding, the proposed method is named Multi-Order BiLSTM (MO-BiLSTM). In this section, we first give the details of the training and the decoding process, and then introduce a pruning technique to improve the efficiency of MO-BiLSTM.

3.1 Multi-Order Training

Our proposed multi-order- n model is a mixture of k single order models with different orders, where n is the maximum order of the single order models. When $n = 1$, the multi-order model becomes a single order-1 model, i.e. a BiLSTM. The order set of the single order models is the subset of $\{1, 2, \dots, n\}$. For example, if the maximum order n is 3, the combination of the single order models can be $[1, 2]$, $[1, 3]$, $[2, 3]$, or $[1, 2, 3]$. Formally, we denote the order set as $\{o_1, o_2, \dots, o_k\}$, where $o_i < o_j$ and $i < j$. In our implementation, n is equal to k in both training and decoding stage.

At the training stage, we train k single order models separately following Eq. 3:

$$\theta_i = \operatorname{argmax}_{\theta} s_{o_i}(y_1, y_2, \dots, y_T | \mathbf{x}; \theta) = \operatorname{argmax}_{\theta} \prod_{t=1}^T s(y_{t-o_i+1} \dots y_t | \mathbf{x}; \theta) \quad (4)$$

where θ_i is the parameters of the i -th single order model of the order o_i . After training, we obtain a set of k independent models: $\{s(y_{t-o_1+1} \dots y_t | \mathbf{x}; \theta_1), \dots, s(y_{t-o_k+1} \dots y_t | \mathbf{x}; \theta_k)\}$, which learns the label dependency of different orders.

3.2 Multi-Order Decoding

For the purpose of simplicity and clarity, we first describe the proposed decoding method of MO-BiLSTM in the order-2 case, and then we extend it to the general order- n case.

Algorithm 1 Multi-order decoding with pruning in the order- n case

- 1: **Input:** sentence x , trained order-1 LSTM $s_1(y|x)$ in Eq. 1, multi-order- n LSTM $s_n(y|x)$ in Eq. 6
 - 2: **for** $t = 1 \dots T$ **do**
 - 3: Select the top- k uni-labels by the order-1 scores:
 - 4: $\tilde{Y}_1 = \text{topkTag}(s_1(y_t|x)), \tilde{Y}_2 = \text{topkTag}(s_1(y_{t-1}|x)), \dots, \tilde{Y}_n = \text{topkTag}(s_1(y_{t-n+1}|x))$
 - 5: Combine n top- k uni-label sets into a n -gram label set:
 - 6: $Y = \tilde{Y}_1 \times \tilde{Y}_2 \times \dots \times \tilde{Y}_n$
 - 7: **for each** $(y^1, y^2, \dots, y^n) \in Y$ **do**
 - 8: Previous tag state $d_{t-1} = y^1 y^2 \dots y^{n-1}$
 - 9: Current tag state $d_t = y^2 y^3 \dots y^n$
 - 10: Compute the transition score $s = s_n(y^1, y^2, \dots, y^n|x)$ by multi-order- n LSTM
 - 11: Compute the maximum score at current state $A[t][d_t] = \max(A[t][d_t], A[t-1][d_{t-1}] * s)$
 - 12: **Output:** The optimal tag sequence y^* by backtracking the path of the maximum score $A[T][d_T]$
-

As shown in Figure 2, in the order-2 case the multi-order model is a mixture of 2 single order models, i.e. single order-1 model (Eq. 1) and single order-2 model (Eq. 2). At the decoding stage, the multi-order model takes account of both the order-1 model and the order-2 model. We need a new decoding approach to unify the decisions of both models. Since the order-1 model and order-2 model predict the label sequence independently, we choose to multiply the scores of order-1 model and order-2 model to get a global score, and use a dynamic programming algorithm to search for the label sequence with the maximum score:

$$\begin{aligned} y_1^*, y_2^*, \dots, y_T^* &= \underset{\mathbf{y}}{\operatorname{argmax}} s_1(y_1, y_2, \dots, y_T | \mathbf{x}; \theta_1) \times s_2(y_1, y_2, \dots, y_T | \mathbf{x}; \theta_2) \\ &= \underset{\mathbf{y}}{\operatorname{argmax}} \prod_{t=1}^T s(y_t | \mathbf{x}; \theta_1) \times s(y_{t-1}, y_t | \mathbf{x}; \theta_2) \end{aligned} \quad (5)$$

where $s(y_t | \mathbf{x}; \theta_1)$ and $s(y_{t-1}, y_t | \mathbf{x}; \theta_2)$ are the score predictions of the single order-1 model and the single order-2 model, respectively. The details of the dynamic programming algorithm are shown in Section 3.3.

Further, we extend the order-2 case to a general order- n case. The difference with the order-2 case is that there are k single order models to approximate the scores of the generated label sequence. We approximate the scores by multiplying all the scores of these trained single order models, and then decode the sequence with the maximum score. Formally, it can be written as:

$$\begin{aligned} y_1^*, y_2^*, \dots, y_T^* &= \underset{\mathbf{y}}{\operatorname{argmax}} \prod_{i=1}^k s_{o_i}(y_1, y_2, \dots, y_T | \mathbf{x}; \theta_i) \\ &= \underset{\mathbf{y}}{\operatorname{argmax}} \prod_{i=1}^k \prod_{t=1}^T s(y_{t-o_i+1} \dots y_t | \mathbf{x}; \theta_i) \end{aligned} \quad (6)$$

where $s(y_{t-o_i+1} \dots y_t | \mathbf{x}; \theta_i)$ is the score prediction of the i -th single order model of the order o_i .

3.3 Scalable Decoding with Pruning

Here, we introduce an efficient dynamic programming algorithm to search for the label sequence with the maximum score. The scores of different n -gram labels are jointly considered in our model. Originally, we should consider all possible n -gram labels at every position of the sentence during dynamic programming. However, it will lead to a huge search space and a lot of time. In order to reduce the time cost, we can prune the unnecessary searching branches. For example, an order-1 model assigns a very low probability to the uni-label “I” of the t -th word, which means the order-1 model is confident that the t -th word can hardly be labeled as “I”. Therefore, it is unnecessary to take account of the bi-gram labels “I-B”, “I-I”, and “I-O” at the next time step.

Model	All-Chunking	English-NER	Dutch-NER
Single Order-1 BiLSTM	93.89	88.23	77.20
Single Order-2 BiLSTM	93.71 (-0.18)	87.61 (-0.62)	76.61 (-0.59)
Single Order-3 BiLSTM	93.34 (-0.55)	87.47 (-0.76)	76.47 (-0.73)
Multi-Order-1 BiLSTM	93.89	88.23	77.20
Multi-Order-2 BiLSTM	94.93 (+1.04)	90.23 (+2.00)	80.95 (+3.75)
Multi-Order-3 BiLSTM	95.01 (+1.12)	90.70 (+2.47)	81.76 (+4.56)

Table 1: Results of single order models and MO-BiLSTM. The number in parentheses means the improvements or reductions compared to the results of order-1 models. All-Chunking denotes All-Phrase-Chunking.

In implementation, we use the order-1 labels with high scores to evaluate whether to prune the high order labels. More precisely, we simply keep the top- k order-1 labels at each position. The order- n labels for a specific position is generated by the top- k labels of n tokens around the position. Suppose a task has totally 50 labels. The order-1 model should compute 50 scores of these labels at each time step. As for the order-3 model, the number of the scores to be computed becomes 50^3 . The original search space before pruning for dynamic programming at each time step is 50^3 . But if we only keep top-5 order-1 labels at each position and prune the order- n labels, the search space will be reduced from 50^3 to 5^3 .

According to our experiments, the pruning technique saves a lot of time in the decoding stage and results in no loss of accuracy, and we find top-5 pruning works the best in order to balance the accuracy and the time cost. Details of the experiments can be found in Section 4. Algorithm 1 shows the detailed process of multi-order decoding with pruning in the order- n case.

4 Experiments

4.1 Datasets

Chunking and named entity recognition are sequence labeling tasks that are sensitive to tag dependencies. The tags inside a segment have internal dependencies. The tags in consecutive segments may have dependencies, too. Thus, we conduct experiments on the chunking and NER tasks to evaluate the proposed method. The test metric is F1-score. The chunking data is from CoNLL-2000 shared task (Sang and Buchholz, 2000), where we need to identify constituent parts of sentences (nouns, verbs, adjectives, etc.). To distinguish it from NP-chunking, it is referred to as the all-phrase chunking. We use the English NER data from the CoNLL-2003 shared task (Sang and Meulder, 2003). There are four types of entities to be recognized: PERSON, LOCATION, ORGANIZATION, and MISC. The other NER dataset is the Dutch-NER dataset from the shared task of CoNLL-2002. The types of entities are the same as the English NER dataset.

4.2 Experimental Details

Our model uses a single layer for the forward and backward LSTMs whose dimensions are set to 200. We use the Adam learning method (Kingma and Ba, 2014) with the default hyper parameters. We set the dropout (Srivastava et al., 2014) rate to 0.5.

Following previous work (Huang et al., 2015), we extract some spelling features and context features. We did not use extra resources, with the exception of using Senna embeddings² in Chunking and English-NER tasks. The embeddings in Dutch-NER tasks are randomly initialized with a size of 50. The code is implemented with the python package *Tensorflow* (Abadi et al., 2016).

4.3 Effect of Multi-Order Setting

For simplicity, the single order model of order- n is denoted as single order- n model and the multi-order model in the order- n case is denoted as multi-order- n model. To verify the effectiveness of MO-BiLSTM,

²Downloaded from <http://ronan.collobert.com/senna/>

Model	All-Chunking	English-NER	Dutch-NER
Order-1	14	10	11
Order-2	154	39	44
Order-3	832	138	158

Table 2: The sizes of tag set of different order.

we conduct comparison experiments of single order models and multi-order models. The results are shown in Table 1. The performance of single order BiLSTM models is getting worse with the growing of the order. An intuitive reason is that the increasing size of tag set raises the difficulty to make a correct tag prediction of a word. Although the performance of single high order models is far from satisfactory, the multi-order models perform well with consistent growth of F1-score on three datasets. In chunking, the MO-BiLSTM at order-3 obtains a 18.3% error reduction compared to BiLSTM. It also performs well in the NER tasks, resulting in a 21.6% and a 20.0% error reductions in English-NER and Dutch-NER compared to BiLSTM baselines, respectively.

The results suggest that high order dependency information is indeed beneficial to the prediction. Furthermore, the adopted multi-order setting makes the learned tag dependency specific to the input words. The reason is that the proposed high order model encodes the tag dependency into a single “output tag”, and model the “output tag” relations using a BiLSTM conditioned on the input words. The tag dependency in previous work is represented by a transition matrix, which cannot capture the relations of tag dependencies with respect to the input words. Moreover, MO-BiLSTM can take advantage of the subtle tag dependencies captured by single-order models and naturally integrate multi-order information to make tag prediction. The decoding process of MO-BiLSTM finds a global optimum tag sequence, which significantly reduces the risk of mistakes.

MO-BiLSTM also results in a growing size of tag set. The sizes of tag set from order-1 model to order-3 model are given in Table 2 respectively. The tag size of the model is beyond a hundred at order-3 case. Although the size of tag set grows as the order of model increases, it is acceptable in such sequence labeling problems compared to the vocabulary size in machine translation which can be over millions.

4.4 Effect of Pruning

The effect of pruning on speeding up the decoding is presented in Table 3. As shown, the pruning technique has shown a great ability to save time with no loss of accuracy. We then give a detailed analysis of the pruning technique. Original search process of dynamic programming considers all possible high order dependencies. However, most low-order tags have been assigned very low probabilities by low-order models and they will form almost impossible high-order tags. Thus, we only keep a small subset of all low-order tags, which makes the possible combinations shrink rapidly so that the cost of dynamic programming is greatly reduced. We also find that the pruned search space has no effect on the performance of the models. We suppose it is almost unlikely that the best tag sequence is out of the pruned search space. Hence, the accuracy is kept to the full extent, as shown in our experiments.

Model	All-Chunking		English-NER		Dutch-NER	
	Time (s)	F1	Time (s)	F1	Time (s)	F1
Multi-Order-2 BiLSTM w/o pruning	31.59	94.93	19.23	90.23	26.60	80.95
Multi-Order-2 BiLSTM	13.64	94.93	13.13	90.23	18.42	80.95
Multi-Order-3 BiLSTM w/o pruning	215.21	95.01	51.78	90.70	69.79	81.76
Multi-Order-3 BiLSTM	44.81	95.01	20.43	90.70	28.66	81.76

Table 3: Effect of pruning on speeding up the decoding.

All-Chunking	F1
SVM classifier (Kudo and Matsumoto, 2001)	93.91
Second order CRF (Sha and Pereira, 2003)	94.30
Second order CRF (McDonald et al., 2005)	94.29
Specialized HMM + voting scheme (Shen and Sarkar, 2005)	94.01
Second order CRF (Sun et al., 2008)	94.34
Conv network tagger (senna) (Collobert et al., 2011)	94.32
CRF-ADF (Sun et al., 2014)	94.52
BiLSTM-CRF (Senna) (Huang et al., 2015)	94.46
Edge-based CRF (Ma and Sun, 2016)	94.80
Encoder-decoder-pointer framework(Zhai et al., 2017)	94.72
BiLSTM (our implementation)	93.89
MO-BiLSTM (this work)	95.01

Table 4: All-Chunking: Comparison with state-of-the-art models.

English-NER	F1
Combination of HMM, Maxent etc. (Florian et al., 2003)	88.76
Semi-supervised model combination (Ando and Zhang, 2005)	89.31
Conv-CRF (Senna + Gazetteer) (Collobert et al., 2011)	89.59
CRF with Lexicon Infused Embeddings (Passos et al., 2014)	90.90
BiLSTM-CRF (Senna) (Huang et al., 2015)	90.10
BiLSTM-CRF (Lample et al., 2016)	90.94
BiLSTM-CNNs-CRF (Ma and Hovy, 2016)	91.21
Iterated Dilated CNNs (Strubell et al., 2017)	90.65
CNN-CNN-LSTM (Shen et al., 2018)	90.89
BiLSTM (our implementation)	88.23
MO-BiLSTM (this work)	90.70

Table 5: English-NER: Comparison with state-of-the-art models.

Dutch-NER	F1
AdaBoost (decision trees) (Carreras et al., 2002)	77.05
Semi-structured resources (Nothman et al., 2013)	78.60
Variant of Seq2Seq (Gillick et al., 2015)	78.08
Character-Level Stacked BiLSTM (Kuru et al., 2016)	79.36
BiLSTM-CRF (Lample et al., 2016)	81.74
Special Decoder + Attention (Martins and Kreutzer, 2017)	80.29
BiLSTM (our implementation)	77.20
MO-BiLSTM (this work)	81.76

Table 6: Dutch-NER: Comparison with state-of-the-art models. Gillick et al. (2015) reported a F1-score of 82.84 in their work, but this result is based on multilingual resources.

4.5 Comparison with State-of-the-art Systems

Table 4 shows the results on all-phrase chunking task compared with previous work. We achieve the state-of-the-art performance in all-phrase chunking. Our model outperforms the popular method BiLSTM-CRF (Huang et al., 2015) by a large margin. Shen and Sarkar (2005) also reported a 95.23 F1-score in their paper. However, this result is based on noun phrase chunking (NP-chunking). All phrase chunking task contains much more tags to predict than NP-chunking, so it is more difficult.

GOLD	The ministry updated port conditions and shipping warnings for the Gulf of Mexico (LOC) , Caribbean and Pacific Coast
BiLSTM	The ministry updated port conditions and shipping warnings for the Gulf (LOC) of Mexico(LOC) , Caribbean and Pacific Coast
MO-BiLSTM	The ministry updated port conditions and shipping warnings for the Gulf of Mexico (LOC) , Caribbean and Pacific Coast.
GOLD	About 200 Burmese students marched briefly from troubled Yangon Institute of Technology (ORG) in northern Rangoon on Friday.
BiLSTM	About 200 Burmese students marched briefly from troubled Yangon (LOC) Institute of Technology (ORG) in northern Rangoon on Friday.
MO-BiLSTM	About 200 Burmese students marched briefly from troubled Yangon Institute of Technology (ORG) in northern Rangoon on Friday.

Table 7: Examples of the predictions of BiLSTM and MO-BiLSTM of order-3.

Table 5 shows the comparison results on the English-NER dataset. Ma and Hovy (2016) reported the best result of English NER. The main architecture of their network is BiLSTM-CRF equipped with a CNN layer to extract character-level representations of words. Our model performs slightly worse than it but outperforms BiLSTM-CRFs reported in other papers (Huang et al., 2015; Lample et al., 2016).

The comparison results on Dutch NER are shown in Table 6. Gillick et al. (2015) keeps the best result of Dutch NER. However, the model is trained on four languages. With the monolingual setting, their model achieves 78.08 on F1 score. Another competitive result is reported in the work of Lample et al. (2016). Their model is a BiLSTM-CRF model with an external LSTM layer to extract character-level representations of words. Our model gets the best score when there is no extra resources.

4.6 Case Study

We observe that MO-BiLSTM mainly helps in two aspects: the prediction of boundaries of a segment and the recognition of long segments. Table 7 shows two cases that MO-BiLSTM model predicts correctly but BiLSTM fails to recognize the entities. In the first case, “Gulf of Mexico” should be recognized as the entity “Location”. BiLSTM recognizes “Gulf” and “Mexico” as locations, but fails to recognize “of” as a part of the entity, so that an entire entity is split. The reason is that BiLSTM model predicts the tag independently, and it predicts “O” as the tag of “of” regardless of the neighboring tags. On the contrary, MO-BiLSTM takes account of the neighboring tags, and works well in this case. Considering that both the left tag and the right tag are labeled “LOC”, the word “of” has a larger probability to be a part of the entity.

The second case contains an entity of type “LOC”. BiLSTM succeeds in recognizing the boundary of the entity but predicts a wrong entity type for the word “Yongon”. Although “Yangon” is a city, it should not be recognized as a location because it is a part of an organization. BiLSTM does not consider the neighboring tag, and makes a wrong prediction, while MO-BiLSTM succeeds in predicting a correct entity by considering the neighboring tag.

4.7 Error Analysis

To better analyze the basic model and the MO-BiLSTM, we investigate the cases that can not be handled well in English-NER dataset, and the result is summarized in Figure 3. All the unrecognized entities are classified into five categories, which are “boundary-1”, “boundary-2”, “boundary-3”, “type”, and “no common words”. “Boundary-1” denotes the cases that the gold entity contains a predicted entity, and “boundary-2” means the gold entity is contained by a prediction. “Boundary-3” represents the case that the gold entity and our prediction overlap. “Type” means a entity’s boundaries are recognized correctly but its entity type is misclassified. When there are no common words between the predicted entity and any gold entity, it is denoted as “no common words”. We count the number of wrongly predicted entities of these different categories, and the result is shown in Figure 3a. The “boundary” error (the

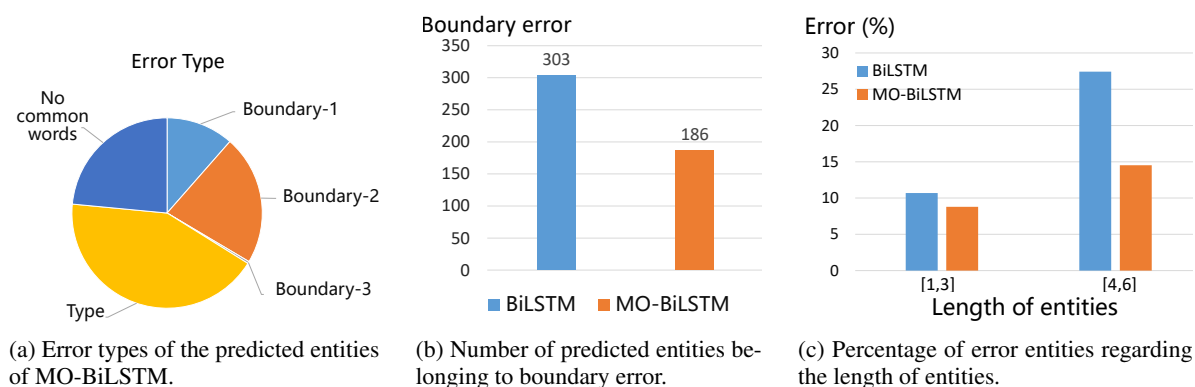


Figure 3: Error analysis of BiLSTM and MO-BiLSTM on English-NER.

sum of “boundary-1”, “boundary-2”, and “boundary-3”), which represents the model misidentifies the entity’s boundaries, is the major error type of BiLSTM. The reason is that the boundary is made up of two tags, but BiLSTM model predicts each tag independently. Our MO-BiLSTM is able to capture the dependencies between two tags, so it can significantly decrease the number of boundary recognition error. That is also the reason why “boundary” error is not the major error of MO-BiLSTM.

We further compare the number of entities belonging to “boundary” error between BiLSTM and MO-BiLSTM. According to Figure 3b, it shows that the “boundary” error of MO-BiLSTM has a reduction rate of nearly 40% compared with BiLSTM. In order to analyze the influence of the length of entities, we divide the entities into 2 groups according to their lengths, and calculate the recognition error rate of different lengths of entities. The result is shown in Figure 3c. We observe that the MO-BiLSTM model has a significant reduction in the recognition error of long entities from 27.42% to 14.52%. The large reduction in error rate proves that the MO-BiLSTM model is able to capture longer distance tag dependencies compared with BiLSTM.

5 Related Work

Huang et al. (2015) and Lample et al. (2016) stacked a CRF layer on BiLSTM to capture the global tag dependencies. The difference between their work is the way to capture character-level information. Their proposed BiLSTM-CRF performs well in sequence labeling tasks. However, the dynamic programming must be done in both training and testing stage. The MO-BiLSTM does not need dynamic programming during training. Muller et al. (2013) proposed a model that also prunes the tag set using a lower order model, but dynamic programming is required in both training and testing stage like prior work. Besides the difference that we do not need dynamic programming in training stage, the pruning technique is different. We directly model the high order states in the training stage, while Muller et al. (2013) merges lower order states to get higher order states. Soltani and Jiang (2016) propose a model called higher order recurrent neural networks (HORNNs). They proposed to use more memory units to keep track of more preceding RNN states, which are all recurrently fed to the hidden layers as feedback. These structures of Soltani’s work are also termed “higher order” models, but the definition is different from ours.

There are several other neural networks that use new techniques to improve sequence labeling. Ling et al. (2015) and Yang et al. (2016) used BiLSTM to compose character embeddings to words representation. Martins and Kreutzer (2017) used an attention mechanism to decide what is the “best” word to focus on next in sequence labeling tasks. Zhai et al. (2017) proposed to separate the segmenting and labeling in chunking. Segmentation is done by a pointer network and a decoder LSTM is used for labeling. Shen et al. (2018) used active learning to strategically choose most useful examples in NER datasets.

6 Conclusions

In this paper, we focus on extending LSTM to higher order models in order to capture more tag dependencies for segmenting and labeling sequence data. We introduce a single order model, which is supposed

to capture more tag dependencies. However, the performance of the single order model is getting worse when increasing the order. To address this problem, we propose to integrate dependency information of different orders to decode. The proposed method, which is called MO-BiLSTM, keeps the scalability to high order models with a pruning technique. Experiments show that MO-BiLSTM achieves better performance than many existing popular methods. It produces the state-of-the-art result in chunking and competitive results in two NER datasets. At the end, we analyze the advantage and limitation of the MO-BiLSTM. We find that MO-BiLSTM mainly helps in the prediction of segment boundaries and the recognition of long segments.

Acknowledgements

This work was supported in part by National Natural Science Foundation of China (No. 61673028), National High Technology Research and Development Program of China (863 Program, No. 2015AA015404), and the National Thousand Young Talents Program. Xu Sun is the corresponding author of this paper.

References

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Gregory S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian J. Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Józefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Gordon Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul A. Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda B. Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2016. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *CoRR*, abs/1603.04467.
- Rie Kubota Ando and Tong Zhang. 2005. A framework for learning predictive structures from multiple tasks and unlabeled data. *Journal of Machine Learning Research*, 6:1817–1853.
- Xavier Carreras, Lluís Marquez, and Lluís Padró. 2002. Named entity extraction using adaboost. In *proceedings of the 6th conference on Natural language learning-Volume 20*, pages 1–4. Association for Computational Linguistics.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel P. Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research*, 12:2493–2537.
- Radu Florian, Abe Ittycheriah, Hongyan Jing, and Tong Zhang. 2003. Named entity recognition through classifier combination. In *Proceedings of the seventh conference on Natural language learning at HLT-NAACL 2003-Volume 4*, pages 168–171. Association for Computational Linguistics.
- Dan Gillick, Cliff Brunk, Oriol Vinyals, and Amarnag Subramanya. 2015. Multilingual language processing from bytes. *CoRR*, abs/1512.00103.
- Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional lstm-crf models for sequence tagging. *arXiv preprint arXiv:1508.01991*.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Taku Kudo and Yuji Matsumoto. 2001. Chunking with support vector machines. In *Proceedings of NAACL'01*, pages 1–8.
- Onur Kuru, Ozan Arkan Can, and Deniz Yuret. 2016. Charner: Character-level named entity recognition. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 911–921.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. *CoRR*, abs/1603.01360.
- Wang Ling, Tiago Luís, Luís Marujo, Ramón Fernández Astudillo, Silvio Amir, Chris Dyer, Alan W. Black, and Isabel Trancoso. 2015. Finding function in form: Compositional character models for open vocabulary word representation. *CoRR*, abs/1508.02096.

- Xuezhe Ma and Eduard H. Hovy. 2016. End-to-end sequence labeling via bi-directional lstm-cnns-crf. *CoRR*, abs/1603.01354.
- Shuming Ma and Xu Sun. 2016. A new recurrent neural CRF for learning non-linear edge features. *CoRR*, abs/1611.04233.
- André FT Martins and Julia Kreutzer. 2017. Learning what’s easy: Fully differentiable neural easy-first taggers. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 349–362.
- Ryan T. McDonald, Koby Crammer, and Fernando Pereira. 2005. Flexible text segmentation with structured multilabel classification. In *HLT/EMNLP 2005, Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference, 6-8 October 2005, Vancouver, British Columbia, Canada*.
- Thomas Muller, Helmut Schmid, and Hinrich Schütze. 2013. Efficient higher-order crfs for morphological tagging. In *EMNLP 2013*.
- Joel Nothman, Nicky Ringland, Will Radford, Tara Murphy, and James R Curran. 2013. Learning multilingual named entity recognition from wikipedia. *Artificial Intelligence*, 194:151–175.
- Alexandre Passos, Vineet Kumar, and Andrew McCallum. 2014. Lexicon infused phrase embeddings for named entity resolution. *arXiv preprint arXiv:1404.5367*.
- Erik Tjong Kim Sang and Sabine Buchholz. 2000. Introduction to the CoNLL-2000 shared task: Chunking. In *Proceedings of CoNLL’00*, pages 127–132.
- E. F. Sang and F. D. Meulder. 2003. Introduction to the CoNLL-2003 Shared Task: Language-Independent Named Entity Recognition. In *Proceedings of CoNLL-2003*, pages 142–147.
- Fei Sha and Fernando C. N. Pereira. 2003. Shallow parsing with conditional random fields. In *HLT-NAACL*.
- Hong Shen and Anoop Sarkar. 2005. Voting between multiple data representations for text chunking. In *Advances in Artificial Intelligence, 18th Conference of the Canadian Society for Computational Studies of Intelligence, Canadian AI 2005, Victoria, Canada, May 9-11, 2005, Proceedings*, pages 389–400.
- Yanyao Shen, Hyokun Yun, Zachary C. Lipton, Yakov Kronrod, and Animashree Anandkumar. 2018. Deep active learning for named entity recognition. In *International Conference on Learning Representations*.
- Rohollah Soltani and Hui Jiang. 2016. Higher order recurrent neural networks. *CoRR*, abs/1605.00064.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15:1929–1958.
- Emma Strubell, Patrick Verga, David Belanger, and Andrew McCallum. 2017. Fast and accurate sequence labeling with iterated dilated convolutions. *arXiv preprint arXiv:1702.02098*.
- Xu Sun, Louis-Philippe Morency, Daisuke Okanohara, and Jun’ichi Tsujii. 2008. Modeling latent-dynamic in shallow parsing: A latent conditional model with improved inference. In *Proceedings of COLING’08*, pages 841–848, Manchester, UK.
- Xu Sun, Wenjie Li, Houfeng Wang, and Qin Lu. 2014. Feature-frequency-adaptive on-line training for fast and accurate natural language processing. *Computational Linguistics*, 40(3):563–586.
- Xu Sun. 2014. Structure regularization for structured prediction. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 2402–2410.
- Zhilin Yang, Ruslan Salakhutdinov, and William W. Cohen. 2016. Multi-task cross-lingual sequence tagging from scratch. *CoRR*, abs/1603.06270.
- Feifei Zhai, Saloni Potdar, Bing Xiang, and Bowen Zhou. 2017. Neural models for sequence chunking. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, pages 3365–3371.