

High-order Refining for End-to-end Chinese Semantic Role Labeling

Hao Fei¹, Yafeng Ren² and Donghong Ji^{1*}

1. Department of Key Laboratory of Aerospace Information Security and Trusted Computing, Ministry of Education, School of Cyber Science and Engineering, Wuhan University, China

2. Guangdong University of Foreign Studies, China

{hao.fe, renyafeng, dhji}@whu.edu.cn

Abstract

Current end-to-end semantic role labeling is mostly accomplished via graph-based neural models. However, these all are first-order models, where each decision for detecting any predicate-argument pair is made in isolation with local features. In this paper, we present a high-order refining mechanism to perform interaction between all predicate-argument pairs. Based on the baseline graph model, our high-order refining module learns higher-order features between all candidate pairs via attention calculation, which are later used to update the original token representations. After several iterations of refinement, the underlying token representations can be enriched with globally interacted features. Our high-order model achieves state-of-the-art results on Chinese SRL data, including CoNLL09 and Universal Proposition Bank, meanwhile relieving the long-range dependency issues.

1 Introduction

Semantic role labeling (SRL), as the shallow semantic parsing aiming to detect the semantic predicates and their argument roles in texts, plays a core role in natural language processing (NLP) community (Pradhan et al., 2005; Zhao et al., 2009; Lei et al., 2015; Xia et al., 2019b). SRL is traditionally handled by two pipeline steps: predicate identification (Scheible, 2010) and argument role labeling (Pradhan et al., 2005). More recently, growing interests are paid for developing end-to-end SRL, achieving both two subtasks, i.e., recognizing all possible predicates together with their arguments jointly, via one single model (He et al., 2018a).

The end-to-end joint architecture can greatly alleviate the error propagation problem, thus helping to achieve better task performance. Currently, the end-to-end SRL methods largely are graph-based

neural models, enumerating all possible predicates and their arguments exhaustively (He et al., 2018a; Cai et al., 2018; Li et al., 2019). However, these first-order models that only consider one predicate-argument pair at a time can be limited to short-term features and local decisions, thus being subjective to long-range dependency issues existing at large surface distances between arguments (Chen et al., 2019; Lyu et al., 2019). This makes it imperative to capture the global interactions between multiple predicates and arguments.

In this paper, based on the graph-based model architecture, we propose to further learn the higher-order interaction between all predicate-argument pairs by performing iterative refining for the underlying token representations. Figure 1 illustrates the overall framework of our method. The BiLSTM encoder (Hochreiter and Schmidhuber, 1997) first encodes the inputs into the initial token representations for producing predicate and argument representations, respectively. The biaffine attention then exhaustively calculates the score representations for all the candidate predicate-argument pairs. Based on all these score representations, our high-order refining module generates high-order feature for each corresponding token via an attention mechanism, which is then used for upgrading the raw token representation. After total N iterations of the above refining procedure, the information between the predicates and the associated arguments can be fully interacted, and thus results in global consistency for SRL.

On the other hand, most of the existing SRL studies focus on the English language, while there is little work in Chinese, mainly due to the limited amount of annotated data. In this study, we focus on the Chinese SRL. We show that our proposed high-order refining mechanism can be especially beneficial for such lower-resource language. Meanwhile, our proposed refining process is fully

*Corresponding author.




Figure 1: The overview of the graph-based high-order model for end-to-end SRL. The dotted-line green box is our proposed high-order refining module.

parallel and differentiable.

We conduct experiments on the dependency-based Chinese SRL datasets, including CoNLL09 (Hajič et al., 2009), and Universal Proposition Bank (Akbik et al., 2015; Akbik and Li, 2016). Results show that the graph-based end-to-end model with our proposed high-order refining consistently brings task improvements, compared with baselines, achieving state-of-the-art results for Chinese end-to-end SRL.

2 Related Work

Gildea and Jurafsky (2000) pioneer the task of semantic role labeling, as a shallow semantic parsing. Earlier efforts are paid for designing hand-crafted discrete features with machine learning classifiers (Pradhan et al., 2005; Punyakanok et al., 2008; Zhao et al., 2009). Later, a great deal of work takes advantages of neural networks with distributed features (FitzGerald et al., 2015; Roth and Lapata, 2016; Marcheggiani and Titov, 2017; Strubell et al., 2018). On the other hand, many previous work shows that integrating syntactic tree structure can greatly facilitate SRL (Marcheggiani et al., 2017; He et al., 2018b; Zhang et al., 2019; Fei et al., 2020b).

Prior studies traditionally separate SRL into two individual subtasks, i.e., predicate disambiguation and argument role labeling, mostly conducting only the argument role labeling based on the pre-

identified predicate (Pradhan et al., 2005; Zhao et al., 2009; FitzGerald et al., 2015; He et al., 2018b; Fei et al., 2020a). More recently, several researches consider the end-to-end solution that handles both two subtasks by one single model. All of them employs graph-based neural model, exhaustively enumerating all the possible predicate and argument mentions, as well as their relations (He et al., 2018a; Cai et al., 2018; Li et al., 2019; Xia et al., 2019a). Most of these end-to-end models, however, are first-order, considering merely one predicate-argument pair at a time. In this work, we propose a high-order refining mechanism to reinforce the graph-based end-to-end method.

Note that most of the existing SRL work focuses on the English language, with less for Chinese, mainly due to the limited amount of annotated data (Xia et al., 2019a). In this paper, we aim to improve the Chinese SRL and make compensation of the data scarcity by our proposed high-order model.

3 Framework

Task formulation. Following prior end-to-end SRL work (He et al., 2018a; Li et al., 2019), we treat the task as predicate-argument-role triplets prediction. Given an input sentence $S = \{w_1, \dots, w_n\}$, the system is expected to output a set of triplets $\mathcal{Y} \in \mathcal{P} \times \mathcal{A} \times \mathcal{R}$, where $\mathcal{P} = \{p_1, \dots, p_m\}$ are all possible predicate tokens, $\mathcal{A} = \{a_1, \dots, a_l\}$ are all associated argument tokens, and \mathcal{R} are the corresponding role labels for each a_i , including a null label ϵ indicating no relation between a pair of predicate argument.

3.1 Baseline Graph-based SRL Model

Our baseline SRL model is mostly from He et al. (2018a). First, we obtain the vector representation x_t^w of each word w_t from pre-trained embeddings. We then make use of the part-of-speech (POS) tag for each word, and use its embedding x_t^{pos} . A convolutional neural networks (CNNs) is used to encode Chinese characters inside a word x_t^c . We concatenate them as input representations: $x_t = [x_t^w; x_t^{pos}; x_t^c]$.

Thereafter, a multi-layer bidirectional LSTM (BiLSTM) is used to encode the input representations into contextualized token representations: $h_1, \dots, h_n = \text{BiLSTM}(x_1, \dots, x_n)$. Based on the token representations, we further generate the separate predicate representations and argument representations: $v_t^p = \text{FFN}(h_t)$, $v_t^a = \text{FFN}(h_t)$.

Then, a biaffine attention (Dozat and Manning, 2016) is used for scoring the semantic relationships exhaustively over all the predicate-argument pairs:

$$\mathbf{v}^s(p_i, a_j) = \mathbf{v}_i^p \cdot \mathbf{W}_1 \cdot \mathbf{v}_j^a + \mathbf{W}_2 \cdot [\mathbf{v}_i^p; \mathbf{v}_j^a] + \mathbf{b}, \quad (1)$$

where \mathbf{W}_1 , \mathbf{W}_2 and \mathbf{b} are parameters.

Decoding and learning. Once a predicate-argument pair (p_i, a_j) (i.e., the role label $r \neq \epsilon$) is determined by a softmax classifier, based on the score representation $\mathbf{v}^s(p_i, a_j)$, the model outputs this tuple (p, a, r) .

During training, we optimize the probability $P_\theta(\hat{y}|S)$ of the tuple $y_{(p_i, a_j, r)}$ over a sentence S :

$$\begin{aligned} P_\theta(y|S) &= \prod_{p \in \mathcal{P}, a \in \mathcal{A}, r \in \mathcal{R}} P_\theta(y_{(p, a, r)}|S) \\ &= \prod_{p \in \mathcal{P}, a \in \mathcal{A}, r \in \mathcal{R}} \frac{\phi(p, a, r)}{\sum_{\hat{r} \in \mathcal{R}} \phi(p, a, \hat{r})}, \end{aligned} \quad (2)$$

where θ is the parameters of the model and $\phi(p, a, r)$ represents the total unary score from:

$$\begin{aligned} \phi(p, a, r) &= \mathbf{W}_p \text{ReLU}(\mathbf{v}^p) + \mathbf{W}_a \text{ReLU}(\mathbf{v}^a) \\ &\quad + \mathbf{W}_s \text{ReLU}(\mathbf{v}^s(p, a)). \end{aligned} \quad (3)$$

The final objective is to minimize the negative log-likelihood of the golden structure:

$$\mathcal{L} = -\log P_\theta(y|S). \quad (4)$$

3.2 Higher-order Refining

The baseline graph model is a first-order model, since it only considers one predicate-argument pair (as in Eq. 3) at a time. This makes it limited to short-term and local decisions, and thus subjective to long-distance dependency problem whenever there are larger surface distances between arguments. We here propose a higher-order refining mechanism for allowing a deep interactions between all predicate-argument pairs.

Our high-order model is shown in Figure 1. Compared with the baseline model, the main difference lies in the high-order refining module. Our motivation is to inform each predicate-argument pair with the information of the other rest of pairs from the global viewpoint. We reach this by refining the underlying token representations \mathbf{h}_t with refined ones which carry high-order interacted features.

Concretely, we take the baseline as the initiation, performing refinement iteratively. At the i -th refining iteration, we can collect the score representations $\mathbf{V}^{i,s} = \{\mathbf{v}_1^{i,s}, \dots, \mathbf{v}_K^{i,s}\}$ of all candidate

predicate-argument pairs, where K (i.e., $\binom{n}{2}$) are the total combination number of these pairs. Based on $\mathbf{V}^{i,s}$, we then generate the high-order feature vector \mathbf{o}_t^i by using an attention mechanism guided by the current token representation \mathbf{h}_t^{i-1} for word w_t at last turn, i.e., the $(i-1)$ -th iteration:

$$\begin{aligned} u_k^i &= \tanh(\mathbf{W}_3 \mathbf{h}_t^{i-1} + \mathbf{W}_4 \mathbf{v}_k^{i,s}), \\ \alpha_k^i &= \text{softmax}(u_k^i), \\ \mathbf{o}_t^i &= \sum_{k=1}^K \alpha_k^i \mathbf{v}_k^{i,s}, \end{aligned} \quad (5)$$

where \mathbf{W}_3 and \mathbf{W}_4 are parameters. We then concatenate the raw token representation and high-order feature representation together, and obtain the refined token representation after a non-linear projection $\hat{\mathbf{h}}_t^i = \text{FFN}([\mathbf{o}_t^i; \mathbf{h}_t^{i-1}])$. Finally, we use $\hat{\mathbf{h}}_t^i$ to update the old one \mathbf{h}_t^{i-1} . After total N iterations of high-order refinement, we expect the model to capture more informative features at global scope and achieve the global consistency.

4 Experiments

4.1 Settings

Our method is evaluated on the Chinese SRL benchmarks, including CoNLL09¹ and Universal Proposition Bank (UPB)². Each dataset comes with its own training, development and test sets. Precision, recall and F1 score are used as the metrics.

We use the pre-trained Chinese fasttext embeddings³. The BiLSTM has hidden size of 350, with three layers. The kernel sizes of CNN are [3,4,5]. We adopt the Adam optimizer with initial learning rate of 1e-5. We train the model by mini-batch size in [16,32] with early-stop strategy. We also use the contextualized Chinese word representations, i.e., ELMo⁴ and BERT (Chinese-base-version)⁵.

4.2 Main Results

We mainly make comparisons with the recent end-to-end SRL models, as well as the pipeline methods on standalone argument role labeling given the gold predicates. Table 1 shows the results on the Chinese CoNLL09. We first find that the joint detection for predicates and arguments can be more beneficial

¹<https://catalog.ldc.upenn.edu/LDC2012T03>

²<https://github.com/System-T/UniversalPropositions>

³<https://fasttext.cc/>

⁴<https://github.com/HIT-SCIR/ELMoForManyLangs>

⁵<https://github.com/google-research/bert>

	Arg.			Prd.
	P	R	F1	F1
• Pipeline method				
Zhao et al. (2009)	80.4	75.2	77.7	-
Björkelund et al. (2009)	82.4	75.1	78.6	-
Roth and Lapata (2016)	83.2	75.9	79.4	-
Marcheggiani and Titov (2017)	84.6	80.4	82.5	-
He et al. (2018b)	84.2	81.5	82.8	-
Cai and Lapata (2019) [‡]	85.4	84.6	85.0	-
• End-to-end method				
He et al. (2018a)	82.6	83.6	83.0	85.7
Cai et al. (2018)	84.7	84.0	84.3	86.0
Li et al. (2019)	84.9	84.6	84.8	86.9
Xia et al. (2019a)	84.6	85.7	85.1	87.2
+BERT	88.0	89.1	88.5	89.6
Ours	85.7	86.2	85.9	88.6
+ELMo	86.4	87.6	87.1	88.9
+BERT	87.4	89.3	88.8	90.3

Table 1: Performances on CoNLL09. Results with [‡] indicates the additional resources are used.

	P	R	F1
He et al. (2018a)	64.8	65.3	64.9
Cai et al. (2018)	65.0	66.4	65.8
Li et al. (2019)	65.4	67.2	66.0
Xia et al. (2019a)	65.2	67.6	66.1
Ours	<u>67.5</u>	<u>68.8</u>	<u>67.9</u>
+ELMo	68.0	70.6	68.8
+BERT	70.0	73.0	72.4

Table 2: Performances by end-to-end models for the argument role labeling on UPB.

than the pipeline detection of SRL, notably with 85.1% F1 score on argument detection by Xia et al. (2019a). Most importantly, our high-order end-to-end model outperforms all these baselines on both two subtasks, with 85.9% F1 score for argument role labeling and 88.6% F1 score for predicate detection. When the contextualized word embeddings are available, we find that our model can achieve further improvements, i.e., 88.8% and 90.3% F1 scores for two subtasks, respectively.

Table 2 shows the performances on UPB. Overall, the similar trends are kept as that on CoNLL09. Our high-order model still performs the best, yielding 67.9% F1 score on argument role labeling, verifying its prominent capability for the SRL task. Also with BERT embeddings, our model further wins a great advance of performances.

4.3 Analysis

High-order refinement. We take a further step, looking into our proposed high-order refining




Figure 2: Performances by varying refining iterations.

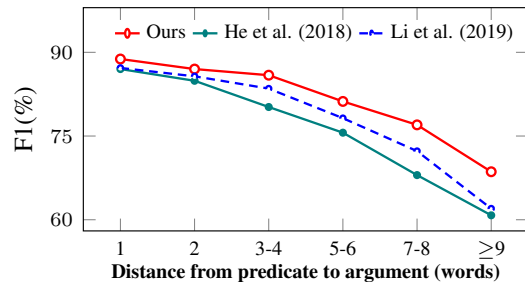


Figure 3: Argument recognition under varying surface distance between predicates and arguments.

mechanism. We examine the performances under varying refining iterations in Figure 2. Compared with the first-order baseline model by He et al. (2018a), our high-order model can achieve better performances for both two subtasks. We find that our model can reach the peak for predicate detection with total 2 iterations of refinement, while the best iteration number is 4 for argument labeling.

Long-distance dependencies. Figure 3 shows the performances of argument recognition by different surface distance between predicates and arguments. The overall results decrease when arguments are farther away from the predicates. Nevertheless, our high-order model can beat against such drop significantly. Especially when the distance grows larger, e.g., distance ≥ 7 , the winning score by our model even becomes more notable.

5 Conclusion

We proposed a high-order end-to-end model for Chinese SRL. Based on the baseline graph-based model, our high-order refining module performed interactive learning between all predicate-argument pairs via attention calculation. The generated higher-order featured token representations then were used to update the original ones. After total N iterations of refinement, we enriched the underlying token representations with global interactions, and made the learnt features more informative. Our

high-order model brought state-of-the-art results on Chinese SRL data, i.e., CoNLL09 and Universal Proposition Bank, meanwhile relieving the long-range dependency issues.

Acknowledgments

We thank the anonymous reviewers for their valuable and detailed comments. This work is supported by the National Natural Science Foundation of China (No. 61772378, No. 61702121), the National Key Research and Development Program of China (No. 2017YFC1200500), the Research Foundation of Ministry of Education of China (No. 18JZD015), the Major Projects of the National Social Science Foundation of China (No. 11&ZD189), the Key Project of State Language Commission of China (No. ZDI135-112) and Guangdong Basic and Applied Basic Research Foundation of China (No. 2020A151501705).

References

- Alan Akbik, Laura Chiticariu, Marina Danilevsky, Yunyao Li, Shivakumar Vaithyanathan, and Huaiyu Zhu. 2015. Generating high quality proposition banks for multilingual semantic role labeling. In *Proceedings of the ACL*, pages 397–407.
- Alan Akbik and Yunyao Li. 2016. Polyglot: Multilingual semantic role labeling with unified labels. In *Proceedings of the ACL*, pages 1–6.
- Anders Björkelund, Love Hafdell, and Pierre Nugues. 2009. Multilingual semantic role labeling. In *Proceedings of the CoNLL*, pages 43–48.
- Jiaxun Cai, Shexia He, Zuchao Li, and Hai Zhao. 2018. A full end-to-end semantic role labeler, syntactic-agnostic over syntactic-aware? In *Proceedings of the CoLing*, pages 2753–2765.
- Rui Cai and Mirella Lapata. 2019. Semi-supervised semantic role labeling with cross-view training. In *Proceedings of the EMNLP*, pages 1018–1027.
- Xinchi Chen, Chunchuan Lyu, and Ivan Titov. 2019. Capturing argument interaction in semantic role labeling with capsule networks. In *Proceedings of the EMNLP*, pages 5415–5425.
- Timothy Dozat and Christopher D Manning. 2016. Deep biaffine attention for neural dependency parsing. *arXiv preprint arXiv:1611.01734*.
- Hao Fei, Meishan Zhang, and Donghong Ji. 2020a. Cross-lingual semantic role labeling with high-quality translated training corpus. In *Proceedings of the ACL*, pages 7014–7026.
- Hao Fei, Meishan Zhang, Fei Li, and Donghong Ji. 2020b. Cross-lingual semantic role labeling with model transfer. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:2427–2437.
- Nicholas FitzGerald, Oscar Täckström, Kuzman Ganchev, and Dipanjan Das. 2015. Semantic role labeling with neural network factors. In *Proceedings of the EMNLP*, pages 960–970.
- Daniel Gildea and Daniel Jurafsky. 2000. Automatic labeling of semantic roles. In *Proceedings of the ACL*, pages 512–520.
- Jan Hajič, Massimiliano Ciaramita, Richard Johansson, Daisuke Kawahara, Maria Antònia Martí, Lluís Màrquez, Adam Meyers, Joakim Nivre, Sebastian Padó, Jan Štěpánek, Pavel Straňák, Mihai Surdeanu, Nianwen Xue, and Yi Zhang. 2009. The CoNLL-2009 shared task: Syntactic and semantic dependencies in multiple languages. In *Proceedings of the CoNLL*, pages 1–18.
- Luheng He, Kenton Lee, Omer Levy, and Luke Zettlemoyer. 2018a. Jointly predicting predicates and arguments in neural semantic role labeling. In *Proceedings of the ACL*, pages 364–369.
- Shexia He, Zuchao Li, Hai Zhao, and Hongxiao Bai. 2018b. Syntax for semantic role labeling, to be, or not to be. In *Proceedings of the ACL*, pages 2061–2071.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9(8):1735–1780.
- Tao Lei, Yuan Zhang, Lluís Màrquez, Alessandro Moschitti, and Regina Barzilay. 2015. High-order low-rank tensors for semantic role labeling. In *Proceedings of the NAACL*, pages 1150–1160.
- Zuchao Li, Shexia He, Hai Zhao, Yiqing Zhang, Zhuosheng Zhang, Xi Zhou, and Xiang Zhou. 2019. Dependency or span, end-to-end uniform semantic role labeling. In *Proceedings of the AAAI*, pages 6730–6737.
- Chunchuan Lyu, Shay B. Cohen, and Ivan Titov. 2019. Semantic role labeling with iterative structure refinement. In *Proceedings of EMNLP*, pages 1071–1082.
- Diego Marcheggiani, Anton Frolov, and Ivan Titov. 2017. A simple and accurate syntax-agnostic neural model for dependency-based semantic role labeling. In *Proceedings of the CoNLL*, pages 411–420.
- Diego Marcheggiani and Ivan Titov. 2017. Encoding sentences with graph convolutional networks for semantic role labeling. In *Proceedings of the EMNLP*, pages 1506–1515.
- Sameer Pradhan, Wayne Ward, Kadri Hacioglu, James Martin, and Daniel Jurafsky. 2005. Semantic role labeling using different syntactic views. In *Proceedings of the ACL*, pages 581–588.

- Vasin Punyakanok, Dan Roth, and Wen-tau Yih. 2008. The importance of syntactic parsing and inference in semantic role labeling. *Computational Linguistics*, 34(2):257–287.
- Michael Roth and Mirella Lapata. 2016. Neural semantic role labeling with dependency path embeddings. In *Proceedings of the ACL*, pages 1192–1202.
- Christian Scheible. 2010. An evaluation of predicate argument clustering using pseudo-disambiguation. In *Proceedings of the LREC*, pages 1187–1194.
- Emma Strubell, Patrick Verga, Daniel Andor, David Weiss, and Andrew McCallum. 2018. Linguistically-informed self-attention for semantic role labeling. In *Proceedings of the EMNLP*, pages 5027–5038.
- Qingrong Xia, Zhenghua Li, and Min Zhang. 2019a. A syntax-aware multi-task learning framework for Chinese semantic role labeling. In *Proceedings of the EMNLP*, pages 5382–5392.
- Qingrong Xia, Zhenghua Li, Min Zhang, Meishan Zhang, Guohong Fu, Rui Wang, and Luo Si. 2019b. Syntax-aware neural semantic role labeling. In *Proceedings of the AAAI*, pages 7305–7313.
- Yue Zhang, Rui Wang, and Luo Si. 2019. Syntax-enhanced self-attention-based semantic role labeling. In *Proceedings of the EMNLP*, pages 616–626.
- Hai Zhao, Wenliang Chen, and Chunyu Kit. 2009. Semantic dependency parsing of NomBank and PropBank: An efficient integrated approach via a large-scale feature selection. In *Proceedings of the EMNLP*, pages 30–39.