

Span Model for Open Information Extraction on Accurate Corpus

Junlang Zhan,^{1,2,3} Hai Zhao^{1,2,3,*}

¹Department of Computer Science and Engineering, Shanghai Jiao Tong University

²Key Laboratory of Shanghai Education Commission for Intelligent Interaction and Cognitive Engineering, Shanghai Jiao Tong University, Shanghai, China

³MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University
longmr.zhan@sjtu.edu.cn, zhaohai@cs.sjtu.edu.cn

Abstract

Open Information Extraction (Open IE) is a challenging task especially due to its brittle data basis. Most of Open IE systems have to be trained on automatically built corpus and evaluated on inaccurate test set. In this work, we first alleviate this difficulty from both sides of training and test sets. For the former, we propose an improved model design to more sufficiently exploit training dataset. For the latter, we present our accurately re-annotated benchmark test set (Re-OIE2016) according to a series of linguistic observation and analysis. Then, we introduce a span model instead of previous adopted sequence labeling formulation for n -ary Open IE. Our newly introduced model achieves new state-of-the-art performance on both benchmark evaluation datasets.

Introduction

Open Information Extraction (Open IE) aims to generate a structured representation of information from natural language text in the form of verbs (or verbal phrases) and their arguments. An Open IE system is usually domain-independent and does not rely on a pre-defined ontology schema. An example of Open IE is shown in Table 1. Open IE has been widely applied in many downstream tasks such as textual entailment and question answering (Fader, Zettlemoyer, and Etzioni 2014). Some existing Open IE systems are listed in Table 2. Most of these systems are built in unsupervised manner except that TextRunner and OLLIE are built in a self-supervised manner.

To perform Open IE in supervised learning and to leverage the advantages of neural network, (Stanovsky et al. 2018) created the first annotated corpus by an automatic translation from Question-Answer Driven Semantic Role Labeling (QA-SRL) annotations (He, Lewis, and Zettlemoyer 2015) and Question-Answer Meaning Representation (QAMR) (Michael et al. 2018). Based on this corpus, they developed an Open IE system RnnOIE using a bidirectional

Sentence	<i>Repeat customers can purchase luxury items at reduced prices.</i>
Extraction	(A0: repeat customers; can purchase ; A1: luxury items; A2: at reduced prices)

Table 1: An example of Open IE in a sentence. The extraction consists of a predicate (in bold) and a list of arguments, separated by semicolons.

long short-term memory (BiLSTM) labeler and BIO tagging scheme which built the first supervised learning model for Open IE. (Cui, Wei, and Zhou 2018) constructed a large but noisy annotated corpus by using OpenIE4 to perform extractions on Wikipedia then kept only tuples with high confidence score. They also built an Open IE system by using a neural sequence to sequence model and copy mechanism to generate extractions. However, this system can only perform binary extractions.

In this work, we develop an Open IE system by adapting a modified span selection model which has been applied in Semantic Role Labeling (SRL) (Ouchi, Shindo, and Matsumoto 2018; He et al. 2018b; Li et al. 2019; Zhao, Zhang, and Kit 2013; Zhao, Chen, and Kit 2009; Li et al. 2018b; Cai et al. 2018), coreference resolution (Shou and Zhao 2012; Zhang, Wu, and Zhao 2012) and syntactic parsing (Zhang, Zhao, and Qin 2016; Ma and Zhao 2012; Li, Zhao, and Parnow 2020; Zhou and Zhao 2019; Zhao et al. 2009; Li et al. 2018a). The advantage of a span model is that span level features can be sufficiently exploited which is hard to perform in token based sequence labeling models. As far as we know, this is the first attempt that applies span model on Open IE.

As previous Open IE systems were mostly built on a loose data basis, we thus intend to strengthen such a basis from both sides of training and test sets. We constructed a large training corpus following the method of (Cui, Wei, and Zhou 2018). The differences of our construction method and the method of (Cui, Wei, and Zhou 2018) are two-fold: First, our corpus is constructed for n -ary extraction instead of binary extraction. Second, previous methods only keep extractions with high confidence score for training. However, the confi-

*Corresponding author. This paper was partially supported by National Key Research and Development Program of China (No. 2017YFB0304100), Key Projects of National Natural Science Foundation of China (U1836222 and 61733011).
Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Systems		Base	Main Features
Name	Work		
TextRunner	(Yates et al. 2007)	-	The first Open IE system using self-supervised learning approach.
ReVerb	(Fader, Soderland, and Etzioni 2011)	-	Leveraging POS tag patterns.
OLLIE	(Mausam et al. 2012)	Reverb	Learning extraction patterns.
ClausIE	(Del Corro and Gemulla 2013)	-	Decomposing a sentence into clauses.
Stanford	(Angeli, Johnson Premkumar, and Manning 2015)	-	Splitting a sentence into utterances.
OpenIE4	(Mausam 2016)	OLLIE	Adding RELNOUN (Pal and 2016), SRLIE (Christensen et al. 2011).
OpenIE5	-	OpenIE4	Adding CALMIE (Saha and 2018), BONIE (Saha, Pal, and Mausam 2017).
PropS	(Stanovsky et al. 2016)	-	Using conversion rules from dependency trees.
MinIE	(Gashteovski, Gemulla, and Del Corro 2017)	ClausIE	Mitigating overspecific extractions.
Graphene	(Cetto et al. 2018)	-	Parsing a sentence into core facts and accompanying contexts.
RnnOIE	(Stanovsky et al. 2018)	-	BiLSTM and BIO tagging.
Seq2seq OIE	(Cui, Wei, and Zhou 2018)	-	Seq2seq model and copy mechanism.

Table 2: A summary of existing Open IE systems

dence scores estimated by OpenIE4 are not so reliable and we have found that some extractions with low confidence score are also useful. We thus propose a method designed to leverage such extractions with low confidence scores.

At the stage of evaluation, the benchmark constructed by (Stanovsky and Dagan 2016) is widely used as the test set for Open IE performance evaluation. However, since this corpus was automatically constructed from QA-SRL, we observe that there are still incorrect extractions. To reduce the noise and obtain a more accurate evaluation for Open IE systems, we manually relabel this corpus. Experiments show that our system outperforms previous Open IE systems in both old benchmark and our relabeled benchmark.

Model

Our system, named SpanOIE, consists of two parts. The first part is the predicate module to find predicate spans in a sentence. The second part is the argument module which takes a sentence and a predicate span as input and outputs argument spans for this predicate.

Problem Definition

We formalize the open information extraction as a span selection task. A span (i, j) is a piece of a sentence beginning with the word i and ending with the word j . The goal is to select the correct span for each Open IE role label. Given a sentence $W = w_1, \dots, w_n$, our model predicts a set of predicate-argument relations $Y \subseteq S \times P \times L$, where $S = \{(w_i, \dots, w_j) | 1 \leq i \leq j \leq n\}$ is the set of all the spans in W , P is the set of predicate spans which is a subset of S and $L = \{A0, A1, A2, A3\}$ is the label set of Open IE¹. For each label l , our model selects a span (i, j) with the highest

¹According to (Stanovsky et al. 2018), this labeled set covers more than 96% sentences.

score:

$$\arg \max_{(i', j') \in S} SCORE_l(i', j'), l \in L \quad (1)$$

where

$$SCORE_l(i, j) = P_\theta(i, j | l) = \frac{\exp(\phi_\theta(i, j, l))}{\sum_{(i', j') \in S} \exp(\phi_\theta(i', j', l))} \quad (2)$$

and ϕ_θ is a trainable scoring function with parameters θ . To train the parameters θ , in the training set, for each sample X and the gold structure Y^* , we minimize the cross-entropy loss:

$$l_\theta(X, Y^*) = \sum_{(i, j, l) \in Y^*} -\log P_\theta(i, j | l) \quad (3)$$

Note that some labels may not appear in the given sentence. In this case, we define the predicate span as NULL span and train a model to select the NULL span.

Spans Candidates Selection

The main drawback of the span model is that too many spans generated from a sentence will cause too high computational cost and hurt the performance of the model. The number of all possible spans for a sentence with size T is $\frac{T(T+1)}{2}$, which makes the number of the spans grow quickly with the length of the sentence. To reduce the number of spans, we propose three constraints for span candidates:

- **Maximum length constraint:** During training, we only keep the argument spans whose size is less than 10 words and the predicate spans whose size is less than 5 words. At inference stage, we remove this constraint.
- **No overlapping constraint:** We keep only the spans which do not overlap with the given predicate span.

- **Syntactic constraint:** This constraint is only for spans whose length is bigger than 1. We keep only the spans which satisfy that a word is either the syntactic parent of another word in the same span, or the parent of this word is in the same span. For the example shown in Figure 1, the span [*luxury items at*] violates this syntactic constraint because the word *at* is not the parent of a word in the same span and its parent, the word *purchase* is not in this span.

Model Architecture

As our model works as two-stage pipeline, there are two modules for predicate and argument modeling which are implemented in the same way as a BiLSTM network shown in Figure 1.

Given an input sample which contains a sentence S and a predicate span P^2 , we extract a feature vector x_i for every word $w_i \in S$:

$$x_i = emb(w_i) \oplus emb(pos(w_i)) \oplus emb(p(w_i)) \oplus emb(dp(w_i)) \quad (4)$$

where \oplus denotes concatenation, $emb(w_i)$ is the word embedding, $emb(pos(w_i))$ is the POS tag embedding, $emb(p(w_i))$ is the embedding of a binary value which indicates whether w_i belongs to predicate span and $emb(dp(w_i))$ is the embedding of dependency relation between w_i and its syntactic parent.

The feature vector x_i is then fed into a BiLSTM network which computes contextualized output features:

$$\vec{h}_i, \overleftarrow{h}_i = BiLSTM(x_i) \quad (5)$$

The forward and the backward output of the BiLSTM are then concatenated together:

$$h_i = \vec{h}_i \oplus \overleftarrow{h}_i \quad (6)$$

We take span features as follows:

$$f_{span}(s_{i:j}) = h_i \oplus h_j \oplus h_i + h_j \oplus h_i - h_j \quad (7)$$

where $s_{i:j}$ is the span which starts from w_i and ends at w_j . h_i and h_j are used as a part of span features in the works of (He et al. 2018a). $h_i + h_j$ and $h_i - h_j$ are used as span features in (Ouchi, Shindo, and Matsumoto 2018). The span features are then fed into a feed forward network and later a Softmax layer to obtain the scores of different labels for each span. For the predicate model, we set binary labels to indicate if the span is a predicate or not. For the argument model, the predicted labels directly correspond to the argument labels.

Inference

At the inference stage, several effective decoding methods have been proposed for the span selection model such as structural constraint inference by using integer linear programming and dynamic programming. In this work, we

²The argument span model will takes the known predicate span as input, while for the predicate span model, it will only take the sentence as input.

use span-consistent greedy search (Ouchi, Shindo, and Matsumoto 2018).

Predicate Inference A predicate span that is completely included in another predicate span in the same sentence will be dropped. For example, given a sentence [*James wants to sell his company*], two spans [*wants to sell*] and [*to sell*] are both selected as predicates, we keep only the former one. This strategy is reasonable since [*James; to sell; his company*] is not an exact fact that can be derived from the original sentence.

Argument Inference Given a sentence and a predicate span, our argument model first scores every possible span for each label in the sentence. Then all the tuples ($span$, $label$, $score$) are sorted according to their scores and top scored tuples are selected. If a tuple with a specific label was not selected before, then the tuple will be selected. The selection ends until all the labels have their corresponding span. If the corresponding span of a label is the NULL span, this label will be dropped in the extraction.

Confidence Score

Like other Open IE systems, our model also provides a confidence score for every extraction. The confidence score CS of an extraction is defined as follows:

$$CS = score_{pred} + \sum_{i=1}^{label_number} score_{e_i, span_i} \quad (8)$$

where $score_{pred}$ is the score for the predicate given by the output of Softmax layer of the predicate model, $score_{e_i, span_i}$ is the score of selected span with label i given by the span model.

Usage of Training Corpus

We use a part of raw corpus that is preprocessed by (Cui, Wei, and Zhou 2018)³. OpenIE4⁴ is used to do the n -ary extraction. Although there exists OpenIE5⁵, which is an advanced version of OpenIE4, we still use the latter due to too high computational cost⁶ required by the former. Actually, The extractions of OpenIE4 and OpenIE5 are the same for most cases and the major improvements from OpenIE5 are about numerical information extraction and splitting conjunctive sentences, which are not a major concern for our preprocessing.

(Cui, Wei, and Zhou 2018) adopt a preprocessing to keep only the extractions with the confidence score greater than 0.9 for training. However, we observe that the confidence scores provided by OpenIE4 are not all reliable. For example, in the sentence [*he makes a state visit*], The extraction by OpenIE4 is [*he; makes; a state visit*]. This extraction looks quite correct while its confidence score is only 0.388 out of 1. We further observe that all the extractions that contain *he*, *she*, *they*, *it* are given low scores. Since Open IE

³<https://1drv.ms/u/s!ApPZxTWwibImHI49ZBwxOU0ktHv>

⁴<https://github.com/allenai/openie-standalone>

⁵<https://github.com/dair-iitd/OpenIE-standalone>

⁶The processing time is much longer in OpenIE5 than in OpenIE4

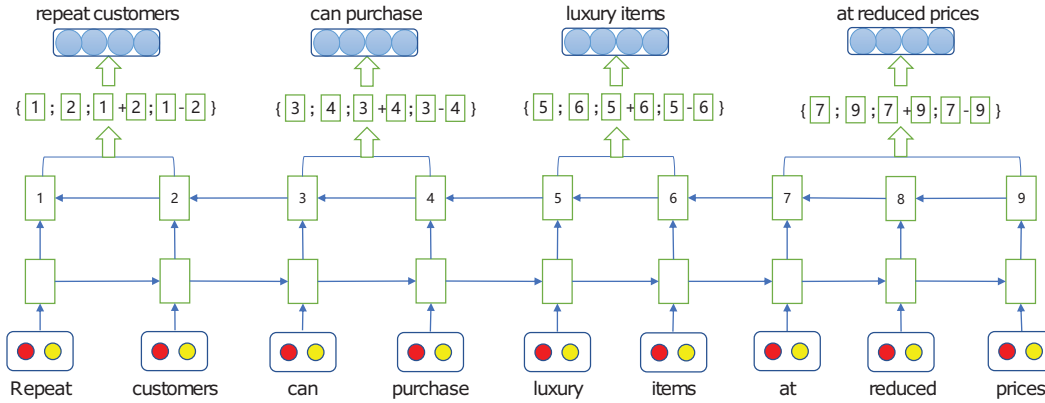


Figure 1: An overview of the span model

surely belongs to the task of information extraction which is supposed to care about non-trivial constituents, and pronoun seems less interesting to the task. However, we argue that extraction with pronoun may sometimes be quite useful. For example, if the sentence [*he makes a state visit*] is in the article about Barack Obama, then *he* is referred to the non-trivial named entity Barack Obama. Ignoring the *he* will cause a loss of critical clue. The pronouns with low scores are just one of many scoring problems caused by OpenIE4 according to our empirical observation. As (Cui, Wei, and Zhou 2018) takes an absolute threshold 0.9 to filter out all extractions below the threshold, we may lose some useful training data. Thus instead of setting a better threshold, we propose to make use of all the extractions from OpenIE4 by adapting our model loss function in Equation (3) to accommodate the confidence scores CS_4 given by OpenIE4,

$$l_{\theta}(X, Y^*) = CS_4(Y^*) \times \sum_{(i,j,l) \in Y^*} -\log P_{\theta}(i, j|l) \quad (9)$$

In detail, we leverage all the extractions obtained by OpenIE4 for training but for every extraction, we multiply the cross entropy by its confidence score according to OpenIE4. In this way, open to all extractions, our span model may automatically determine the selection of spans, for either those correct extractions but with low confidence scores or those truly bad extractions with high confidence scores.

In this training corpus, there are 1,109,411 sentences and 2,175,294 extractions. In addition, we annotated this corpus with POS tag and dependency parsing information by using spaCy⁷. The positions of the first word and the last word for every span are provided.

Test Corpus: Relabeled OIE2016

OIE2016 is a broadly adopted evaluation benchmark corpus for open IE, which was used as test set created by (Stanovsky and Dagan 2016)⁸. This corpus which consists of 600 sen-

tences⁹ is the processing results of QA-SRL (He, Lewis, and Zettlemoyer 2015).

However, we have observed some problematic annotations and extractions in OIE2016:

1. OIE2016 considers all the verbal adjectives as predicates, which causes overspecification of extractions. For the sentence in Figure 1, OIE2016 gives an extraction [*reduced; prices*]. In the construction of QA-SRL corpus, annotators were asked to find all possible arguments for every verb, which causes such an extraction overspecification.
2. OIE2016 considers appositive as A0 while ignoring the appositive relations. For the sentence [*Now Mr. Broberg, a lawyer, claims he'd play for free*], OIE2016 gives two extractions, [*Mr. Broberg; claims; that he'd play for free*] and [*a lawyer; claims; that he'd play for free*], in which the latter is less meaningful. In addition, a meaningful extraction [*Mr. Broberg; he; a lawyer*] is missing.
3. There are numbers of incorrect extractions due to the wrong translation from QA-SRL. For the sentence [*A cafeteria is located on the sixth floor, a chapel on the 14th floor, and a study hall on the 15th floor*], OIE2016 gives two incorrect extractions, [*a study hall; located; on the sixth floor*] and [*a study hall; located; on the 14th floor*].

To fix all the above problems, we manually re-annotated the entire OIE2016 corpus for a more accurate evaluation benchmark. The re-annotated OIE2016, named Re-OIE2016, has solved the problem of extraction overspecification, and provides extractions for appositive relation. Besides the predicates and arguments, we also annotated context information in Re-OIE2016. However, since some Open IE systems (including ours) do not perform the context extraction as it is hard to judge the correct boundary for context, this information does not serve as a part of evaluation in this work and will be exploited in future research.

⁹Originally, there are 3200 sentences in OIE2016. In this paper, we only consider a subset of this corpus. This subset is also used as test data in (Stanovsky et al. 2018).

⁷<https://spacy.io>

⁸<https://github.com/gabrielStanovsky/oie-benchmark>

Model Parameters	
Word embedding size	100
POS tag embedding size	10
Dependency label size	20
Hidden size of BiLSTM	200
Number of layers of BiLSTM	2
Training Parameters	
Batch size	20
Training steps	55470
Optimizer	Adam
Learning rate	0.01
Dropout	0.3
Decay rate for every 100 steps	0.005

Table 3: Experiment settings

Experiments¹⁰

Experiment Settings

In the experiments, we use pre-trained GloVe¹¹ as word embedding. Other experiment settings are shown in Table 3.

Open IE Systems

We will evaluate the performance of different Open IE systems on both OIE2016 and Re-OIE2016. These systems include ClausIE, OLLIE, OpenIE4, Props, Stanford OIE and our system SpanOIE. To demonstrate the advantage of the span model compared to the sequence labeling model, we also realized a sequence labeling model which adapts BIO tagging scheme. The training process and the parameters settings of the BIO model are the same as SpanOIE except that the output of the BiLSTM layer is directly fed into a feed forward network to do the classification for BIO labels. To focus on the model own strength, our model will not include dependency embeddings in the input representation unless stated otherwise.

Results on OIE2016

We use OIE2016 to evaluate the precision and recall of different Open IE systems. The Precision-Recall (P-R) curve is shown in Figure 2. The P-R curve is obtained by selecting a list of confidence score thresholds for each Open IE system. All the extractions below this threshold are dropped. We see that the SpanOIE system outperforms all other systems. Although our model is trained from the corpus bootstrapped from the extractions of OpenIE4, it still has a better performance than OpenIE4 as it is capable of using truly good extractions no matter how the threshold varies. The improvement of our model over OpenIE4 is two-fold. First, our model can find more predicates than OpenIE4 which leads to higher recall on more extractions. Second, our model is more accurate in finding the correct arguments. As an example shown in Table 4, our model finds one more predicate [*to access*] and the correct A0 span [*The keys*].

¹⁰Our code and annotated corpus are publicly available in https://github.com/zhanjunlang/Span_OIE

¹¹<https://nlp.stanford.edu/projects/glove/>




Figure 2: The P-R curve of different Open IE systems on OIE2016




Figure 3: The P-R curve of different Open IE systems on Re-OIE2016

Results on Re-OIE2016

We also evaluate the Open IE systems on Re-OIE2016 corpus. The P-R curves are shown in Figure 3. We also compare the Area under P-R curve (AUC) and the best F1 score for different Open IE systems on two corpora in Table 5. Note that the code for the work of (Cui, Wei, and Zhou 2018) is not released and it is hard to reproduce their result. So we just report their results obtained from their paper.

From the comparison, we see that although the performance ranking of Open IE systems is kept the same between OIE2016 and Re-OIE2016, the AUC value and the best F1 score increase for all the Open IE systems on Re-OIE2016. Comparing P-R curves, we see that the increase of AUC value mainly comes from the increase of recall. In fact, the number of extractions is fewer in Re-OIE2016 than OIE2016. The decrease of extractions mainly comes from three aspects that have been mentioned in previous sections. First, OIE2016 considers all the verbal adjectives as predicates which are deleted in Re-OIE2016. Second, it also excludes the extractions that take appositive as A0.

Third, OIE2016 breaks conjunctions in arguments to generate multiple extractions. For example, given a sentence [*Ben and Jenny went to Japan*], if the extracting is done

Original sentence	<i>The keys, which were needed to access the building, were locked in the car.</i>
SpanOIE	[<i>The keys; were needed; to access the building</i>] [<i>The keys; to access; the building</i>] [The keys ; were locked; in the car]
OpenIE4	[<i>The keys; were needed; to access the building</i>] [<i>The keys, which were needed; were locked; in the car</i>]

Table 4: Example of extractions from Span OIE and OpenIE 4

Systems	OIE2016		Re-OIE2016	
	AUC	F1	AUC	F1
Stanford (Angeli, Johnson Premkumar, and Manning 2015)	0.079	13.55	0.115	16.70
OLLIE (Mausam et al. 2012)	0.202	38.58	0.313	49.47
PropS (Stanovsky et al. 2016)	0.320	54.38	0.433	64.16
ClausIE (Del Corro and Gemulla 2013)	0.364	58.01	0.464	64.17
OpenIE4 (Mausam 2016)	0.408	58.83	0.509	68.32
BIO	0.462	68.55	0.579	77.08
Seq2seq OIE (Cui, Wei, and Zhou 2018)	0.473	-	-	-
SpanOIE_prev	0.491	69.42	0.647	78.08
SpanOIE	0.489	68.65	0.659	78.50

Table 5: AUC and the highest F1 score for different Open IE systems for OIE2016 and Re-OIE2016. SpanOIE_prev is the model trained from the extractions with confidence score higher than 0.9 of OpenIE4. The rest settings are the same as SpanOIE.

on both sides of the conjunction, we have [*Ben; went to; Japan*] and [*Jenny; went to; Japan*] for extractions. However, since most existing Open IE systems (including ours) do not break the conjunctions, we preserve the conjunction constituent in an extraction instead of breaking them. Actually, the conjunction structures can be better postprocessed after the works of Open IE. Note that there are also numbers of wrong extractions in OIE2016 due to the incorrect break of conjunctions. Thus the extractions over the conjunctions are not taken into account in our Re-OIE2016 for evaluating Open IE systems.

Usage of Training Corpus

During the training, we make use of every extraction from OpenIE4 in the training corpus by integrating the confidence score into the loss function. Following (Cui, Wei, and Zhou 2018) which keeps only the extractions with confidence score higher than 0.9, we compare our span model to every Open IE system on both OIE2016 and Re-OIE2016. From Table 5, we can see that the performance of two approaches is very close, which indicates that the use of extractions with low confidence score does not hurt much the performance.

We find that our method helps improve the model in some cases. For example, for the sentence [*And he was in Ali 's army in the Battle of Jamal and later it was Muhammad ibn Abu Bakr who escorted Aisha back to Madina*] in OIE2016, for the predicate *was*, the extraction from previous model is [*Muhammad ibn Abu Bakr; was; in Ali 's army*] and the extraction from our model is [*he; was; in Ali 's army; in the Battle of Jamal*]. From this example, we see that our model is capable of performing pronoun argument extracting while the previous model cannot due to the absence of training

corpus with pronoun argument.

Besides the extraction for pronoun arguments, we also observe that our model has improvements on other extraction patterns. For example, for the sentence [*Yost was named manager of the Kansas City Royals , replacing Trey Hillman*] with the predicate *replacing*, the extraction from the previous model is [*the Kansas City Royals; replacing; Trey Hillman*] while the extraction from our model is [*Yost; replacing; Trey Hillman*]. This is an extraction pattern where the predicate is in the form of present participle while the subject is far from the predicate.

The Impact of Syntax

To explore if syntactic clues may help Open IE, we introduce syntactic dependency labels into the word level representation of our model. The dependency parsing is performed by spaCy. Based on the span feature designed in Equation (7), we add another vector in this feature:

$$f_{span}(s_{i:j}) = h_i \oplus h_j \oplus h_i + h_j \oplus h_i - h_j \oplus h_{dp(s_{i:j})} \quad (10)$$

where $dp(s_{i:j})$ is the index of the parent word of syntactic head of the span $s_{i:j}$. For example, Figure 4 illustrates syntactic dependency parse tree for the span [*luxury items*], which is a part of the sentence in Figure 1. The syntactic head of this span is the word *items*, whose parent is the word *purchase*. So h_4 , which is the BiLSTM output for the word *purchase*, is also taken as a part of the span features in this example. Syntactic parse tree discloses the nonlinear structure for a span and is supposed to make a help in Open IE because all the words are not equally important in the same span. For example, in the span [*luxury items*], the syntactic clue indicates that the word *items* is more important than the

word *luxury* in finding the correct label for this span. The latter one might be more helpful in determining the boundary of the span.




Figure 4: Syntactic dependency tree for [*luxury items*]

Using the syntax-aware span feature as Equation (10), we re-train our model and achieve AUC 0.487 and 0.639 on OIE2016 and Re-OIE2016, respectively, which are both slightly lower than our original syntax-agnostic model. Such experimental results undoubtedly strike the belief on the helpfulness of syntactic information, which may be attributed to the noise inside the training corpus and imperfect predicted syntactic parsing trees.

Despite the unfavorable performance, we still found that the span level syntactic information helps improve the extraction in some cases. For example, given a sentence [*The dialects they speak are similar but have different intonations*], with the predicate *have*, syntax-agnostic model gives the extraction [*they; have; different intonations*] but the extraction from a syntax-aware model is [*The dialects; have; different intonations*]. This example suggests that syntax is indeed helpful in Open IE, and it is hopefully proved that an effective syntax-aware Open IE system can be built on training corpus of higher quality and more accurate syntactic inputs.

Limitations

This subsection discusses cases on limitations of the current Open IE systems, including ours.

Open IE systems like OLLIE, OpenIE4 and ClausIE also extract context of a sentence while most of the latest neural model cannot do so. For example, given a sentence [*He believes The Lakers will win the game*], the correct extraction should be [*context[He believes]; The Lakers; will win; the game*] because [*The Lakers will win the game*] is not true according to the original sentence. That is also an important difference between Open IE and SRL since Open IE needs to identify correct and useful information. The reason why we did not bootstrap context information from OpenIE4 is that the context information extracted by OpenIE 4 is very noisy. Second, most of neural Open IE systems including ours cannot extract is appositive relation. For example, given a sentence [*Obama, president of USA, gave a speech on Friday*], one extraction should be [*Obama; be; president of USA*].

Third, we still need to explore better ways in constructing and leveraging training corpus for Open IE. One possible way is that we can bootstrap tuples from different existing Open IE systems instead of just from OpenIE4 so as to leverage the strength of different systems.

Conclusions

In this work, we bootstrap a large training corpus for Open IE and propose an improved model design to more sufficiently exploit the training data. We also present our accurately re-annotated benchmark test set (Re-OIE2016). At last, we introduce SpanOIE, which is the first span model for n -ary Open IE. Our model achieve new state-of-the-art performance on both benchmark evaluation datasets.

References

- Angeli, G.; Johnson Premkumar, M. J.; and Manning, C. D. 2015. Leveraging linguistic structure for open domain information extraction. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (ACL-IJCNLP)*, 344–354.
- Cai, J.; He, S.; Li, Z.; and Zhao, H. 2018. A full end-to-end semantic role labeler, syntax-agnostic over syntax-aware? In *Proceedings of the 27th International Conference on Computational Linguistics (COLING)*.
- Cetto, M.; Niklaus, C.; Freitas, A.; and Handschuh, S. 2018. Graphene: a context-preserving open information extraction system. In *Proceedings of the 27th International Conference on Computational Linguistics (COLING)*, 94–98.
- Christensen, J.; Soderland, S.; Etzioni, O.; et al. 2011. An analysis of open information extraction based on semantic role labeling. In *Proceedings of the sixth international conference on Knowledge capture*, 113–120.
- Cui, L.; Wei, F.; and Zhou, M. 2018. Neural open information extraction. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL) (Volume 2: Short Papers)*, 407–413.
- Del Corro, L., and Gemulla, R. 2013. Clausie: clause-based open information extraction. In *Proceedings of the 22nd international conference on World Wide Web*, 355–366. ACM.
- Fader, A.; Soderland, S.; and Etzioni, O. 2011. Identifying relations for open information extraction. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1535–1545.
- Fader, A.; Zettlemoyer, L. S.; and Etzioni, O. 2014. Open question answering over curated and extracted knowledge bases. In *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.
- Gashteovski, K.; Gemulla, R.; and Del Corro, L. 2017. Minie: Minimizing facts in open information extraction. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2630–2640.
- He, L.; Lee, K.; Levy, O.; and Zettlemoyer, L. 2018a. Jointly predicting predicates and arguments in neural semantic role labeling. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL) (Volume 2: Short Papers)*, 364–369.
- He, S.; Li, Z.; Zhao, H.; and Bai, H. 2018b. Syntax for semantic role labeling, to be, or not to be. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL) (Volume 1: Long Papers)*, 2061–2071.

- He, L.; Lewis, M.; and Zettlemoyer, L. 2015. Question-answer driven semantic role labeling: Using natural language to annotate natural language. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 643–653.
- Li, Z.; Cai, J.; He, S.; and Zhao, H. 2018a. Seq2seq dependency parsing. In *Proceedings of the 27th International Conference on Computational Linguistics (COLING)*, 3203–3214.
- Li, Z.; He, S.; Cai, J.; Zhang, Z.; Zhao, H.; Liu, G.; Li, L.; and Si, L. 2018b. A unified syntax-aware framework for semantic role labeling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2401–2411.
- Li, Z.; He, S.; Zhao, H.; Zhang, Y.; Zhang, Z.; Zhou, X.; and Zhou, X. 2019. Dependency or span, end-to-end uniform semantic role labeling. In *Proceedings of the 33-th AAAI Conference on Artificial Intelligence (AAAI)*.
- Li, Z.; Zhao, H.; and Parnow, K. 2020. Global greedy dependency parsing. In *Proceedings of the 34-th AAAI Conference on Artificial Intelligence (AAAI)*.
- Ma, X., and Zhao, H. 2012. Fourth-order dependency parsing. In *Proceedings of the 2012 International Conference on Computational Linguistics (COLING)*, 785–796.
- Mausam; Schmitz, M.; Soderland, S.; Bart, R.; and Etzioni, O. 2012. Open language learning for information extraction. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-IJCNLP)*, 523–534.
- Mausam, M. 2016. Open information extraction systems and downstream applications. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI)*, 4074–4077.
- Michael, J.; Stanovsky, G.; He, L.; Dagan, I.; and Zettlemoyer, L. 2018. Crowdsourcing question-answer meaning representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 560–568.
- Ouchi, H.; Shindo, H.; and Matsumoto, Y. 2018. A span selection model for semantic role labeling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1630–1642.
- Pal, H., and , M. 2016. Demonyms and compound relational nouns in nominal Open IE. In *Proceedings of the 5th Workshop on Automated Knowledge Base Construction*, 35–39.
- Saha, S., and , M. 2018. Open information extraction from conjunctive sentences. In *Proceedings of the 27th International Conference on Computational Linguistics (COLING)*, 2288–2299.
- Saha, S.; Pal, H.; and Mausam. 2017. Bootstrapping for numerical Open IE. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL)*, 317–323.
- Shou, H., and Zhao, H. 2012. System paper for CoNLL-2012 shared task: Hybrid rule-based algorithm for coreference resolution. In *Joint Conference on EMNLP and CoNLL - Shared Task*, 118–121.
- Stanovsky, G., and Dagan, I. 2016. Creating a large benchmark for open information extraction. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2300–2305.
- Stanovsky, G.; Fidler, J.; Dagan, I.; and Goldberg, Y. 2016. Getting more out of syntax with props. *arXiv preprint arXiv:1603.01648*.
- Stanovsky, G.; Michael, J.; Zettlemoyer, L.; and Dagan, I. 2018. Supervised open information extraction. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT), Volume 1 (Long Papers)*, 885–895.
- Yates, A.; Banko, M.; Broadhead, M.; Cafarella, M.; Etzioni, O.; and Soderland, S. 2007. Texrunner: Open information extraction on the web. In *Proceedings of Human Language Technologies: The Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT)*, 25–26.
- Zhang, X.; Wu, C.; and Zhao, H. 2012. Chinese coreference resolution via ordered filtering. In *Joint Conference on EMNLP and CoNLL - Shared Task*, 95–99.
- Zhang, Z.; Zhao, H.; and Qin, L. 2016. Probabilistic graph-based dependency parsing with convolutional neural network. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL) (Volume 1: Long Papers)*, 1382–1392.
- Zhao, H.; Song, Y.; Kit, C.; and Zhou, G. 2009. Cross language dependency parsing using a bilingual lexicon. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP (ACL-IJCNLP)*, 55–63.
- Zhao, H.; Chen, W.; and Kit, C. 2009. Semantic dependency parsing of NomBank and PropBank: An efficient integrated approach via a large-scale feature selection. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 30–39.
- Zhao, H.; Zhang, X.; and Kit, C. 2013. Integrative semantic dependency parsing via efficient large-scale feature selection. *Journal of Artificial Intelligence Research* 46:203–233.
- Zhou, J., and Zhao, H. 2019. Head-driven phrase structure grammar parsing on Penn treebank. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2396–2408.