

# Integrating Relation Constraints with Neural Relation Extractors

Yuan Ye, Yansong Feng,\* Bingfeng Luo, Yuxuan Lai, Dongyan Zhao

Wangxuan Institute of Computer Technology, Peking University, China

{pkuyeyuan, fengyansong, bf\_luo, erutan, zhaodongyan}@pku.edu.cn

## Abstract

Recent years have seen rapid progress in identifying predefined relationship between entity pairs using neural networks (NNs). However, such models often make predictions for each entity pair individually, thus often fail to solve the inconsistency among different predictions, which can be characterized by discrete relation constraints. These constraints are often defined over combinations of entity-relation-entity triples, since there often lack of explicitly well-defined type and cardinality requirements for the relations. In this paper, we propose a unified framework to integrate relation constraints with NNs by introducing a new loss term, *Constraint Loss*. Particularly, we develop two efficient methods to capture how well the local predictions from multiple instance pairs satisfy the relation constraints. Experiments on both English and Chinese datasets show that our approach can help NNs learn from discrete relation constraints to reduce inconsistency among local predictions, and outperform popular neural relation extraction (NRE) models even enhanced with extra post-processing. Our source code and datasets will be released at <https://github.com/PKUYYeYuan/Constraint-Loss-AAAI-2020>.

## Introduction

Relation extraction (RE) aims to extract predefined relations between two marked entities in plain texts, and its success can benefit many knowledge base (KB) related tasks like knowledge base population (KBP) (Suchanek et al. 2013; Wu, He, and Hu 2018), question answering (QA) (Dai, Li, and Xu 2016; Yu et al. 2017; Lai et al. 2019) and etc.

Most existing works investigate the RE task in a classification style. A sentence marked with a given pair of entities is fed to a classifier to decide their relationship, also called the *sentence-level* RE. Another related setup is to feed a group of sentences containing the given entity pair to the classifier, called the *bag-level* RE. We should note that both *sentence-level* RE and *bag-level* RE make predictions for each entity pair individually and locally. However, when we look at the model outputs globally, there are always contradictions among different predictions, such as an entity is regarded as the object of both *Country* and *City*, two different cities

have been labeled as *Capital* for one country and so on. To alleviate these local contradictions, Chen et al. 2018 collect constraints on the type and cardinality requirements of relations, such as whether two relations should not have the same type of subject (object), or whether a relation should not have multiple subjects (objects) given its object (subject). Further, in the inference stage, they use integer linear programming (*ILP*) to filter and adjust the local predictions that are inconsistent with these constraints. Basically, *ILP* operates in a post-processing way to copy with contradictory predictions, but there is no way to provide feedback to the original RE model.

In fact, it would be of great importance to utilize those constraints to backwards improve the original RE models. For example, enhanced with various attention or pooling mechanisms, most current neural network extraction models have shown promising performance on benchmark datasets, but they still suffer from inconsistent local predictions (Chen et al. 2018). If those relation constraints can be learned during model training, that will help to further improve the overall performance, and we may no longer need a separate post-processing step as *ILP* does.

However, directly integrating relation constraints with NRE models is not a trivial task: (1) relation constraints are not defined regarding a single prediction, but often over combinations of instances, thus it is not easy to find appropriate representations for those constraints; (2) it is not easy to evaluate how well pairwise predictions match the constraints in a batch, and it is not clear how to feed the information back to the NRE models.

To tackle the challenges, we propose a unified framework to flexibly integrate relation constraints with NRE models by introducing a loss term *Constraint Loss*. Concretely, we develop two methods denoted as *Coherent* and *Semantic* to construct *Constraint Loss* from different perspectives. *Coherent* captures how well pairwise predictions match the constraints from an overall perspective, and *Semantic* pays more attention to which specific rule in the constraints the pairwise predictions should satisfy. In addition, we encode relation constraints into different representations for each method. Notably, *Constraint Loss* is regarded as a rule-based regularization term within a batch instead of regularizing

each instance, since the relation constraints are often defined over combinations of local predictions. Moreover, our approach does not bring extra cost to the inference phase and can be adapted to most existing NRE models without explicit modifications to their structures, as it only utilizes the outputs from the NRE model as well as relation constraints to obtain *Constraint Loss* and provides feedback to the NRE model through backward propagation. Experiments on both Chinese and English datasets show that our approach can help popular NRE models learn from the constraints and outperforms state-of-the-art methods even enhanced with *ILP* post-processing. Moreover, jointing our approach and *ILP* achieves further improvement which demonstrates that our approach and the *ILP* post-processing exploit complementary aspects from the constraints.

The main contributions of this paper include: (1) We propose a unified framework to effectively integrate NRE models with relation constraints without interfering the inherent NRE structure. (2) We develop two efficient methods to capture the inconsistency between local NRE outputs and relation constraints, which are used as a loss term to help the NRE training. (3) We provide thoroughly experimental study on different datasets and base models. The results show that our approach is effective and exploits the constraints from different perspectives with *ILP*.

## Related Work

Since annotating high-quality relational facts in sentences is laborious and time-consuming, RE is usually investigated in the distant supervision (DS) paradigm, where datasets are automatically constructed by aligning existing KB triples  $\langle subj, r, obj \rangle$ <sup>1</sup> with a large text corpus (Mintz et al. 2009). However, the automatically constructed dataset suffers the wrong labeling problem, where the sentence that mentions the two target entities may not express the relation they hold in KB, thus contains many false positive labels (Riedel, Yao, and McCallum 2010). To alleviate the wrong labeling problem, RE is usually investigated in the multi-instance learning (MIL) framework which considers RE task at bag-level and holds the at-least-one hypothesis, thinking that there exists at least one sentence which expresses the entity pair’s relation in its corresponding sentence bag (Hoffmann et al. 2011; Surdeanu et al. 2012; Suchanek et al. 2013).

As neural networks have been widely used, an increasing number of researches for RE have been proposed under MIL framework. Zeng et al. 2014 use a convolution neural network (CNN) to automatically extract features and Zeng et al. 2015 use a piece-wise convolution neural network (PCNN) to capture structural information by inherently splitting a sentence into three segments according to the two target entities. Furthermore, Lin et al. 2016 proposed sentence-level attention-based models (ACNN, APCNN) to dynamically reduce the weights of noisy sentences. And there also exists many NN based works improving the RE performance by utilizing external information, such as syntactic information

<sup>1</sup>We use *subj*, *obj* and *r* to denote subject, object and relation for a KB triple, respectively, in the rest of this paper.

(He et al. 2018), entity description (Ji et al. 2017), relation aliases (Vashishth et al. 2018) and etc.

In addition, there are many works focusing on combining NNs with precise logic rules to harness flexibility and reduce uninterpretability of the neural models. Hu et al. 2016 make use of first-order logic (FOL) to express the constraints and propose a teacher-student network that could project prediction probability into a rule-regularized subspace and transfer the information of logic rules into the weights of neural models. Xu et al. 2018 put forward a semantic loss framework, which bridges between neural output vectors and logical constraints by evaluating how close the neural network is to satisfying the constraints on its output with a loss term. And Luo et al. 2018 develop novel methods to exploit the rich expressiveness of regular expressions at different levels within a NN, showing that the combination significantly enhances the learning effectiveness when a small number of training examples are available.

However, using these frameworks on RE is not straightforward. Specifically, Hu et al. 2016 directly project prediction probability of instance as they can assess how well a single instance’s prediction satisfies the rules, while constraints in RE are non-local and we could not examine each instance individually for the violation of constraints. Luo et al. 2018 need the regular expressions to provide keyword information and get a priori category prediction, however, generating high-quality regular expressions from RE datasets is not easy. For Xu et al. 2018, since our constraints are related to the combination of instances rather than a single instance, to utilize the semantic loss framework, we need to find appropriate representations for various relation constraints and evaluate the neural output in a pairwise way.

## Relation Constraints

Since many KBs do not have a well-defined typing system and explicit argument cardinality requirements, similar in spirit with Chen et al. 2018, our relation constraints are defined over the combination of two triples:  $\langle subj_m, r_m, obj_m \rangle$  and  $\langle subj_n, r_n, obj_n \rangle$ .<sup>2</sup> We derive the type and cardinality constraints from existing KB to implicitly capture the expected type and cardinality requirements on the arguments of a relation. One can surely employ human annotators to collect such constraints.

**Type Constraints.** Type constraints implicitly express the types of subjects and objects that a specific relation could have. For example, the subject and object types for relation *almaMater* should be PERSON and SCHOOL, respectively, and we take positive rules [*almaMater* and *knownFor* could have the same subject type] and [*almaMater* and *employer* could have the same object type] to implicitly encode *almaMater*’s subject and object type requirements.

Specifically, we use entity sharing between different relations to implicitly capture the expected argument type of

<sup>2</sup>The main difference is that our constraints are considered as positive rules where we expect the relation predictions to fall in, while the constraints in Chen et al. 2018 are considered as inviolate rules that the local predictions should not break.

each relation. If the subject (or object) set of relation  $r_i$  in KB has an intersection with those of  $r_j$ , then we consider  $r_i$  and  $r_j$  could have the same expected subject (or object) type. We thereby assign relation pairs  $(r_i, r_j)$  into  $C^{ts}$  if they are expected to have the same subject type, into  $C^{to}$  if they are expected to have the same object type, and assign it into  $C^{tso}$  if the subject type of one relation is expected to be same as the object type of the other.

**Cardinality Constraints.** Cardinality constraints indicate the cardinality requirements on a relation’s arguments. For example, relation *almaMater* could have multiple subjects (PERSON) when its object (SCHOOL) is given.

Specifically, for each predefined relation  $r_i$ , we collect all triples containing  $r_i$ , and count the number of the triples that have multiple objects (subjects) for each subject (object). Then, we assign relation  $r_i$  into  $C^{cs}$  if it can have multiple subjects for a given object, into  $C^{co}$  if it can have multiple objects for a given subject.

Finally, we get 5 sub-category constraint sets. We use  $C^\phi$  to represent a single set,  $C^{l*} \in \{C^{ts}, C^{to}, C^{tso}\}$  to represent a type constraint set, and  $C^{c*} \in \{C^{cs}, C^{co}\}$  to represent a cardinality constraint set. Note that our relation constraints are defined to examine whether a pair of subject-relation-object triples can hold at the same time from different perspectives. To make our constraints clearer, we list some rules for each constraint set in Table 1.

Set	Sampled Positive Rules
$C^{ts}$	( <i>almaMater</i> , <i>knowFor</i> ), ( <i>city</i> , <i>region</i> ), ( <i>spouse</i> , <i>child</i> )
$C^{to}$	( <i>almaMater</i> , <i>owner</i> ), ( <i>city</i> , <i>hometown</i> ), ( <i>capital</i> , <i>city</i> )
$C^{tso}$	( <i>birthPlace</i> , <i>capital</i> ), ( <i>child</i> , <i>spouse</i> ), ( <i>city</i> , <i>country</i> )
$C^{cs}$	<i>almaMater</i> , <i>country</i> , <i>city</i> , <i>hometown</i>
$C^{co}$	<i>foundationPerson</i> , <i>child</i> , <i>knownFor</i> , <i>product</i>

Table 1: Example rules for each constraint set  $C^\phi$ .

## Our Approach

As shown in Fig. 1, our framework consists of two main components, a base NRE model and the Constraint Loss Calculator (CLC). The CLC module is designed to integrate the relation constraints with NRE models, which does not rely on specific NRE architectures and can work in a plug-and-play fashion.

### Base NRE Model

While our framework can work with most existing relation extractors, in this paper, we take the most popular neural relation extractors, *ACNN* and *APCNN* (Lin et al. 2016), as our base extractors.<sup>3</sup>

*ACNN* uses convolution neural networks with max-pooling layer to capture the most significant features from

<sup>3</sup>We do not use the most recently neural models, such as Feng et al. 2018, Qin, Xu, and Wang 2018 and Jia et al. 2019, as our base model, as they focused more on noise reduction which is not within the scope of this paper.




Figure 1: Framework overview. For each mini-batch, the *Constraint Loss* is calculated by evaluating the predicted probability  $p_\theta(Y|X)$  according to the relation constraints.

a sentence. Then, an attention layer is used to selectively aggregate individual representations from a bag of sentences into a sentence bag embedding, which is fed to a softmax classifier to predict the relation distribution  $p_\theta(Y|X)$ .

*APCNN* is an extension of *ACNN*. Specifically, *APCNN* divides the convolution output into three segments based on the positions of the two given entities and devises a piecewise max-pooling layer to produce sentence representation.

### Constraint Loss Calculator (CLC)

Given the inherent nature of our relation constraints, the *CLC* can not evaluate a single subject-relation-object prediction against our constraint sets, we thus operate our *CLC* in a mini-batch wise way. Specifically, in each batch, we integrate the relation constraints with the NRE output by introducing a loss term, *Constraint Loss*, which is designed to regulate the NRE model to learn from those constraints, **e.g., not to violate the positive rules in the constraints**. As shown in Fig. 1, to calculate *Constraint Loss*, we first collect the NRE output probability  $p_\theta(Y|X)$  within the batch, and then the *CLC* takes  $p_\theta(Y|X)$  and relation constraints as input to obtain *Constraint Loss*, which should reflect the inconsistency among all local predictions  $p_\theta(Y|X)$  according to our relation constraints. Finally, the total loss for back propagation consists of two parts: the original NRE loss ( $L_O$ ) and the *Constraint Loss* ( $L_C$ ):

$$L_{total} = L_O + \lambda L_C$$

where  $\lambda$  is a weight coefficient.

Particularly, the key task of *CLC* is to evaluate how well the current NRE output probabilities  $p_\theta(Y|X)$  satisfy our relation constraints. We solve this problem in two steps. We first calculate a local loss term  $L(p^m, p^n, C^\phi)$  for a pair of local predictions, i.e.,  $p^m$  and  $p^n$  for the  $m^{th}$  and  $n^{th}$  instances, respectively,<sup>4</sup> against the constraint set  $C^\phi$ . Secondly, we aggregate all local loss terms to obtain the batch-wise *Constraint Loss*. Here, we develop two methods to cal-

<sup>4</sup>We use  $p^m$  to represent the probability output of the neural model on the  $m^{th}$  instance in a batch.

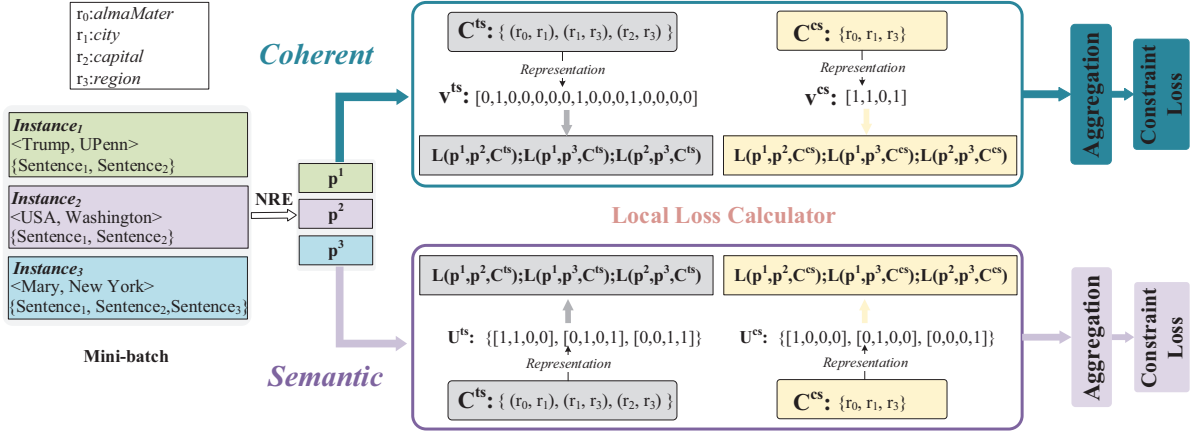


Figure 2: A running example of our two CLC modules, *Semancit* and *Coherent*. To exhibit the main process clearly, we simplify the example and only consider 4 relations, a mini-batch with 3 instances and two constraint sets ( $C^{ts}$ ,  $C^{cs}$ ). The whole process of CLC contains 3 steps, we take *Semantic* as an example. First, we represent constraint set  $C^{ts}$  ( $C^{cs}$ ) as a vector set  $U^{ts}$  ( $U^{cs}$ ) and each vector represents a single rule. Then, we feed NRE output and the vector set into local loss calculator, getting the local loss  $L(p^m, p^n, C^\phi)$  (using Eq. 4) for each pair of instances within a batch. Finally, Constraint Loss is obtained by aggregating all instance pairs in the batch. The main difference between *Semantic* and *Coherent* is that *Coherent* represents constraint set into **one vector** while *Semantic* represents it into a **vector set**, utilizing relation constraints from different perspectives.

culate *Constraint Loss* from different perspectives, denoted as *Coherent (Coh)* and *Semantic (Sem)*, respectively.

### Coherent (Coh)

In this method, we calculate *Constraint Loss* by evaluating the coherence between the NRE output and a constraint set. Note that this method only requires the NRE outputs to be more consistent with one constraint set as a whole, but does not explicitly push the NRE model to update according to specific positive rules in this set.

**Representing Constraint Sets.** We encode a constraint set into **one** single binary vector. Since the positive rules in the type and cardinality constraint set have different forms, we represent them in slightly different ways.

For a type constraint set  $C^{t*}$ , we construct a binary vector  $v^{t*}$ , where  $v_{i,j}^{t*}$  indicates whether relation pair  $(r_i, r_j)$  belongs to  $C^{t*}$ , i.e.,  $v_{i,j}^{t*} = 1$  if  $(r_i, r_j) \in C^{t*}$  and  $v_{i,j}^{t*} = 0$  if  $(r_i, r_j) \notin C^{t*}$ . Take  $C^{ts}$  illustrated in Fig. 2 as an example, since  $(almaMater, city) \in C^{ts}$ ,  $v_{0,1}^{ts}$  is set to 1.

For a cardinality constraint set  $C^{c*}$ , we construct a binary vector  $v^{c*}$ , where  $v_i^{c*}$  indicates whether relation  $r_i$  belongs to  $C^{c*}$ , i.e.,  $v_i^{c*} = 1$  if  $r_i \in C^{c*}$  and  $v_i^{c*} = 0$  if  $r_i \notin C^{c*}$ . Again, in Fig.2,  $v_0^{cs}$  is set to 1, since  $almaMater \in C^{cs}$ .

Thus, for each one of the 5 sub-category constraint sets, we build one single representation vector, resulting in 5 vectors  $\{v^{ts}, v^{to}, v^{tso}, v^{cs}, v^{co}\}$ . And the dimensions of  $v^{t*}$  and  $v^{c*}$  are  $|R|^2$  and  $|R|$ , respectively, where  $|R|$  is the size of the relation set.

**Local Loss Calculation.** Now, we proceed to calculate the loss term for a pair of local predictions, e.g., the  $m^{th}$  and

$n^{th}$  instances, within a batch. Our expectation is that coherent local prediction pairs should satisfy our constraint sets. Again, we deal with the type constraint sets and cardinality constraint sets separately.

Thus, for a type constraint set  $C^{t*}$  represented by  $v^{t*}$ , the local loss,  $L(p^m, p^n, C^{t*})$ , can be written as:

$$L(p^m, p^n, C^{t*}) = -I_{t*}^{mn} \log \left( \sum_{i,j} v_{i,j}^{t*} p_i^m p_j^n \right) \quad (1)$$

where  $I_{t*}^{mn} \in \{0, 1\}$  indicates whether to calculate  $L(p^m, p^n, C^{t*})$ . Take  $I_{ts}^{mn}$  as an example, for triple pair  $(subj_m, r_m, obj_m)$  and  $(subj_n, r_n, obj_n)$ , we set  $I_{ts}^{mn} = 1$ , if  $subj_m = subj_n$  which means the two triples have the same subject type and corresponding predicted relation pair should satisfy  $C^{ts}$ ; otherwise, we assign 0 to  $I_{ts}^{mn}$ .<sup>5</sup>  $p_i^m p_j^n$  can be considered as the probability that the base NRE model predicts relation  $r_i$  and  $r_j$  for the  $m^{th}$  and  $n^{th}$  instances, respectively.

For cardinality constraint set  $C^{c*}$  represented by  $v^{c*}$ , the local loss,  $L(p^m, p^n, C^{c*})$ , can be written as:

$$L(p^m, p^n, C^{c*}) = -I_{c*}^{mn} \log \left( \sum_i v_i^{c*} p_i^m p_i^n \right) \quad (2)$$

where  $I_{c*}^{mn}$  is an indicator similar to  $I_{t*}^{mn}$  and  $p_i^m p_i^n$  is seen as the possibility that the base NRE predicts relation  $r_i$  for both the  $m^{th}$  and  $n^{th}$  instances.

**Aggregation.** To obtain the batch-wise *Constraint Loss*, we simply sum all the local loss terms  $L(p^m, p^n, C^\phi)$  in a batch to get the total constraint loss  $L_C$  (Eq. 3).

$$L_C = \sum_m \sum_n \sum_\phi L(p^m, p^n, C^\phi) \quad (3)$$

<sup>5</sup>Detailed assignment for  $I_\phi^{mn}$  can be found in Appendix.

## Semantic (Sem)

In this method, we pay more attention to which specific rules in the constraint sets the pairwise local predictions should satisfy. Our intuition is that, for each of our constraint set, good local predictions should follow one rule in that set, while bad ones may not find any rules to satisfy. This may push the NRE model to effectively learn from specific rules in a more focused way.

**Representing Constraints.** To represent the rules in the constraint sets more precisely, we encode each rule  $c \in C^\phi$  into a single binary vector  $\mathbf{u}$ , thus, the whole set is represented as a vector set  $U^\phi$ , shown as in Fig. 2. Again, since the rules in  $C^{t*}$  and  $C^{c*}$  have different forms, we represent them in different ways.

For each type rule  $(r_i, r_j) \in C^{t*}$ , the representation vector  $\mathbf{u}$  is a binary vector whose  $i^{th}$  and  $j^{th}$  dimensions are set to 1 and the rest are set to 0. Take  $C^{ts}$  in Fig. 2 as an example, the rule  $(almaMater, city) \in C^{ts}$  is encoded as a vector whose first two dimensions are set to 1.

For each cardinality rule  $r_i \in C^{c*}$ , the representation vector  $\mathbf{u}$  is a binary vector whose  $i^{th}$  dimension is set to 1 and the rest are set to 0. In Fig. 2, the rule  $almaMater \in C^{cs}$  is represented by a vector, where only the first dimension is set to 1.

Different from *Coherent*, here we construct one vector set to represent each sub-category constraint sets, resulting in 5 vector sets  $\{U^{ts}, U^{to}, U^{tso}, U^{cs}, U^{co}\}$ . And each single rule is represented by a  $|R|$ -dim binary vector.

**Local Loss Calculation.** Inspired by the semantic loss function (*SL*) introduced in Xu et al. 2018, which operates on a single output, we adapt the original *SL* to deal with pairwise instances over different kinds of constraints. We design the new local loss term as:

$$L(\mathbf{p}^m, \mathbf{p}^n, C^\phi) = -I_\phi^{mn} \log \sum_{c \in C^\phi} f(\mathbf{p}^m, \mathbf{p}^n, c) \quad (4)$$

where  $I_\phi^{mn}$  is an indicator same as before and  $f$  is a score function reflecting how well the pairwise predictions match a single rule  $c \in C^\phi$ . Since the rules in  $C^{t*}$  and  $C^{c*}$  are encoded in different ways, we calculate  $f$  for type constraint sets and cardinality constraint sets separately.

Thus, for a rule  $c$  in type constraint set  $C^{t*}$ , the score function  $f(\mathbf{p}^m, \mathbf{p}^n, c)$  can be calculated by:

$$q_i = p_i^m + p_i^n - p_i^m p_i^n$$

$$f(\mathbf{p}^m, \mathbf{p}^n, c) = \prod_{u_i=1} q_i \prod_{u_i=0} (1 - q_i) \quad (5)$$

where  $\mathbf{u}$  is the vector representation of  $c$  and  $q_i$  is the probability that base NRE model predicts relation  $r_i$  for at least one of the  $m^{th}$  and the  $n^{th}$  instances.

For a rule  $c$  in cardinality constraint set  $C^{c*}$ ,  $f(\mathbf{p}^m, \mathbf{p}^n, c)$  can be calculated by:

$$f(\mathbf{p}^m, \mathbf{p}^n, c) = \prod_{u_i=1} p_i^m p_i^n \prod_{u_i=0} (1 - p_i^m p_i^n) \quad (6)$$

where  $p_i^m p_i^n$  means the probability that NRE model predicts relation  $r_i$  for both the  $m^{th}$  and the  $n^{th}$  instances.

**Aggregation.** We use the same method as *Coherent* to perform aggregation according to Eq. 3.

Note that *Coherent* handles the constraint set as a whole and treats each single rule in that set equally, while *Semantic* treats all rules in a constraint set as mutually exclusive and makes the pairwise predictions more satisfying one certain rule in that set. Take  $C^{cs}$  as an example, in Eq. 2, *Coherent* just simply increases the probabilities of corresponding relation pairs for all positive rules, and each rule has the same influence on the summation. However, in Eq. 6, for a potentially satisfied rule, *Semantic* not only tries to increase the probabilities of its corresponding relation pair, but also lowers the probabilities of the rest. That is, there would not exist pair-wise local predictions which satisfy two positive rules well in one constraint set at the same time, since if the high probabilities of a relation pair have the positive effect on one specific rule, it has negative effect on all the others.

## Experiments

Our experiments are designed to answer the following questions: (1) whether our approach can effectively utilize the relation constraints to improve the extraction performance? (2) which *CLC* module performs better, *Coherent* or *Semantic*? (3) which is the better way to utilize the relation constraints, learning or post-processing?

### Datasets

We evaluate our approach on both English and Chinese datasets constructed by Chen et al. 2018. The English one is constructed by mapping triples in DBpedia (Bizer et al. 2009) to sentences in the New York Times Corpus. It has 51 relations, about 50k triples, 134k sentences for training and 30k triples, 53k sentences for testing. The Chinese dataset is built by mapping the triples of HudongBaiKe, a large Chinese encyclopedia, with four Chinese economic newspapers. It contains 28 relations, about 60k triples, 120k sentences for training and 40k triples, 83k sentences for testing.

We automatically collect relation constraints for English and Chinese datasets based on corresponding KBs. In total, we obtain 541 rules for the English dataset and 110 rules for the Chinese one.

Here we do not use the popular RE dataset created by Riedel, Yao, and McCallum 2010, since it is produced with an earlier version of Freebase which is not available now, and makes it impossible to automatically collect the constraints. Secondly, Riedel’s dataset is dominated by three *big* relations: *location/contains*, */people/nationality* and */people/place\_lived*, covering about 60% of all KB triples. Therefore, there are not enough data related to other relations for us to collect constraints.

### Setup

Following common practice in the RE community (Ji et al. 2017; He et al. 2018), we report the model performance by both precision-recall (PR) curve and Precision@N (P@N). We also report the average score of P@N (*Mean*).




Figure 3: The PR curves of our approach on two datasets with *ACNN* and *APCNN* as base models.

The main goal of our work is to explore whether our approach can help neural models effectively learn from discrete relation constraints. Therefore, the first baseline models are the two most popular base NRE models, *ACNN* and *APCNN*. We also compare with the base NRE models enhanced with a post-processing *ILP* step, *ACNN+ILP* and *APCNN+ILP*, which can be considered as state-of-the-art constraint-based RE solutions.

We use a grid search to tune our hyper parameters, including the weight coefficient  $\lambda$ . Details about our hyper parameters are reported in Appendix.

## Main Results

Our main results are summarized in Fig. 3 and Table 2. As shown in Fig. 3, we can see that both the red and green dot lines are lifted above the solid black lines, showing that after equipped with our *CLC* modules, i.e., *Coherent* and *Semantic*, both *ACNN* and *APCNN* obtain significant improvement on the English and Chinese datasets. This indicates our *CLC* module actually helps the base NRE models benefit from properly utilizing the relation constraints, without interference to the base models.

However, we find that our approach obtains different levels of improvement on the two datasets. On the Chinese one, as shown in Table 2, with our *Semantic* version *CLC*, *APCNN(Sem)* gains 4.9% improvement in *Mean* compared to *APCNN*, but, on the English dataset, it only receives 0.5% in *Mean*. Similar trends are also found for the *Coherent* version and the *ACNN* base model. The better performance gain on the Chinese dataset is mainly because its relation definitions are more clear compared to that of the English dataset. For example, in English dataset, there are 8 relations whose object could be any *LOCATION*, such as *birthPlace*, while only 3 similar relations exist in Chinese dataset.

In addition, we investigate the performance improvement when applying our *CLC* module to different base NRE models. Although both *ACNN* and *APCNN* are improved by our *CLC* module in various datasets, we can still observe that *ACNN* generally receives more performance improvement compared with the *APCNN* base model. Taking the *Semantic* method as an example, as shown in Table 2, on the English dataset, *ACNN(Sem)* obtains 2.2% performance improvement in *Mean* against *ACNN*, while *APCNN(Sem)* only fetches 0.5% improvement. And similar trends can be found in the *Coherent* method and on the Chinese dataset. The more improvement when taking the *ACNN* as base NRE model is because, compared with *ACNN*, *APCNN* itself is

designed to take the entity-aware sentence structure information into account, thus can extract more effective features that, to some extent, can implicitly capture part of the arguments' type and cardinality requirements of a relation, leaving relatively less space for our *CLC* module to improve.

**Comparing *Coherent* and *Semantic*** This paper presents two different methods, *Coh* and *Sem*, to represent and integrate the relation constraints, both of which can lead to substantial improvement with both base models and datasets. Specifically, as shown in Table 2, *Sem* brings slightly more improvement than *Coh* in most of the settings, e.g., on Chinese dataset, *APCNN(Sem)* obtain about 0.2% more gains (4.9% vs 4.7%) in *Mean* than *APCNN(Coh)*. We think the reason is that *Sem* provides a more precise treatment for the constraints, e.g., embedding each rule with a vector and trying to evaluate the NRE output against one specific rule, while *Coh* represents all rules in a sub-category with one single vector and evaluates the output against whole set of rules, which is admittedly a more coarse fashion.

**Learning? or Post-processing?** Previous works show that *ILP* can effectively solve the inconsistency among predictions in a post-processing fashion (Chen et al. 2018).

Now we discuss which is the better way to utilize the relation constraints, our *CLC* module or traditional *ILP* post-processing. As shown in Table 2, both *APCNN(Sem)* and *APCNN(Coh)* outperforms *APCNN+ILP* by at least 0.1% on the English dataset and 1.0% on the Chinese dataset. Similar trends can be also found for *ACNN(Sem)* and *ACNN(Coh)*. This shows that helping base NRE models to learn from the relation constraints can generally bring more improvement, thus utilizes the constraints more effectively compared to utilizing those constraints in a post-processing way.

We can also apply *ILP* as a post-processing step to our approach, since our *CLC* module works in the model training phase, and leaves the testing phase as it is. Interestingly, as shown in Table 2, with an extra *ILP* post-processing, both *Coh* and *Sem* obtain further improvement with different base NRE models on different datasets. This indicates that our *CLC* module still may not fully exploit the useful information behind the relation constraints. The reasons may be that our approach and the *ILP* post-processing exploit the relation constraints from different perspectives. For example, our *CLC* operates in a mini-batch level during training, that is a relatively local view, but *ILP* post-processing directly optimizes the model output in a slightly global view.

Moreover, in Table 3, we find that applying *ILP* to our

Model Name	English Dataset					Chinese Dataset				
	P@100	P@200	P@300	Mean	$\Delta_{Base}$	P@100	P@200	P@300	Mean	$\Delta_{Base}$
<i>ACNN</i>	96.70	92.61	91.72	93.68	–	89.08	86.89	84.52	86.83	–
<i>ACNN(Coh)</i>	97.39	93.78	90.69	93.96	+0.3	95.86	94.86	93.04	94.59	+7.8
<i>ACNN(Sem)</i>	97.62	95.87	94.12	95.87	+2.2	95.97	94.61	93.53	94.70	+8.1
<i>ACNN+ILP</i>	97.87	94.36	93.16	95.13	+1.5	93.75	92.18	90.10	92.01	+5.2
<i>ACNN(Coh)+ILP</i>	97.73	94.51	91.29	94.51	+0.8	97.09	96.18	94.01	95.76	+9.0
<i>ACNN(Sem)+ILP</i>	<b>98.17</b>	<b>96.6</b>	<b>95.48</b>	<b>96.75</b>	<b>+3.1</b>	<b>97.73</b>	<b>96.40</b>	<b>94.43</b>	<b>96.18</b>	<b>+9.4</b>
<i>APCNN</i>	100	98.97	97.41	98.79	–	92.96	91.75	91.08	91.93	–
<i>APCNN(Coh)</i>	100	99.57	97.33	98.97	+0.2	98.88	96.00	94.98	96.62	+4.7
<i>APCNN(Sem)</i>	100	100	97.95	99.32	+0.5	100	96.97	93.42	96.80	+4.9
<i>APCNN+ILP</i>	100	99.13	97.55	98.89	+0.1	96.06	95.15	94.63	95.28	+3.4
<i>APCNN(Coh)+ILP</i>	100	100	98.03	99.34	+0.6	99.07	96.17	<b>95.16</b>	96.79	+4.9
<i>APCNN(Sem)+ILP</i>	<b>100</b>	<b>100</b>	<b>98.39</b>	<b>99.46</b>	<b>+0.7</b>	<b>100</b>	<b>97.67</b>	94.25	<b>97.31</b>	<b>+5.4</b>

Table 2: Summary P@N(%) scores of our approach on two datasets with *ACNN* and *APCNN* as base models.  $\Delta_{Base}$  indicates the difference between mentioned model and the base NRE model (*ACNN* in the top and *APCNN* in the bottom). And the name with *+ILP* means that we perform *ILP* over the model’s outputs as an extra post-processing.

*CLC* enhanced model receives relatively less gain compared to applying *ILP* to the base model, e.g., 0.5% for *APCNN(Sem)* v.s. 3.4% for *APCNN* on the Chinese dataset. This observation may indicate that our approach has pushed NRE base models to learn part of the useful information behind relation constraints, leaving fewer inconsistent outputs for *ILP* post-processing to filter out. On the other hand, this observation shows again that our *CLC* approach and the *ILP* post-processing exploit complementary aspects from the relation constraints, and our *CLC* module could be further improved by taking more global optimization into account.

	English		Chinese	
	Mean	$\Delta_{ILP}$	Mean	$\Delta_{ILP}$
<i>ACNN</i>	93.68	+1.5	86.83	+5.2
<i>ACNN(Coh)</i>	93.86	+0.6	94.59	+1.2
<i>ACNN(Sem)</i>	<b>95.87</b>	+0.9	<b>94.70</b>	+1.5
<i>APCNN</i>	98.79	+0.1	91.93	+3.4
<i>APCNN(Coh)</i>	98.97	+0.4	96.62	+0.2
<i>APCNN(Sem)</i>	<b>99.32</b>	+0.1	<b>96.80</b>	+0.5

Table 3: Relative improvement of different models in *Mean*.  $\Delta_{ILP}$  is the performance difference between the mentioned model and the same model with an extra *ILP* step. For example,  $\Delta_{ILP}$  corresponding to raw *ACNN* indicates that applying *ILP* to *ACNN* obtains 1.5% and 5.2% gain in *Mean* on English and Chinese dataset, respectively.

## More Analysis

To better understand what our approach learns from the constraints, we take a deep look at the outputs of *APCNN* and *APCNN(Sem)* on the test split of the Chinese dataset. First, we count the total number of contradictory pairwise predictions and find that applying our *Semantic* method to *APCNN*

achieves a reduction of 5,966 violations, 28.0% of the total<sup>6</sup>. This indicates our approach has pushed the base NRE models to learn from the relation constraints. However, there are still many remaining violations since our approach operates during training in a soft and local way, compared to *ILP* during testing.

Another observation is that our approach actually reduces the violations related to each relations, and especially does better when there are tighter requirements on the relation’s arguments. For example, *APCNN(Sem)* reduces 89.6% violations for relation *locationState* compared to *APCNN*, but for *locationRegion*, it only reduces 36.3%. This is because the relation constraints may indicate more clear arguments’ type requirements for *locationState* than those of *locationRegion*, which are captured by our *CLC* module to push into the base NRE during training.

## Conclusion

In this paper, we propose a unified framework to effectively integrate discrete relation constraints with neural networks for relation extraction. Specifically, we develop two approaches to evaluate how well NRE predictions satisfy our relation constraints in a batch-wise, from both general and precise perspectives. We explore our approach on English and Chinese dataset, and the experimental results show that our approach can help the base NRE models to effectively learn from the discrete relation constraints, and outperform popular NRE models as well as their *ILP* enhanced versions. Our study reveals that learning with the constraints can better utilize the constraints from a different perspective compared to the *ILP* post-processing method.

## Acknowledgments

We thank anonymous reviewers for their valuable suggestions. This work is supported in part by the National Hi-Tech

<sup>6</sup>Detailed numbers per category are reported in Appendix.

R&D Program of China (2018YFC0831900) and the NSFC Grants (No.61672057, 61672058). For any correspondence, please contact Yansong Feng.

## References

- Bizer, C.; Lehmann, J.; Kobilarov, G.; Auer, S.; Becker, C.; Cyganiak, R.; and Hellmann, S. 2009. Dbpedia-a crystallization point for the web of data. *Web Semantics: science, services and agents on the world wide web* 7(3):154–165.
- Chen, L.; Feng, Y.; Huang, S.; Luo, B.; and Zhao, D. 2018. Encoding implicit relation requirements for relation extraction: A joint inference approach. *Artificial Intelligence* 265:45–66.
- Dai, Z.; Li, L.; and Xu, W. 2016. Cfo: Conditional focused neural question answering with large-scale knowledge bases. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, 800–810.
- Feng, J.; Huang, M.; Zhao, L.; Yang, Y.; and Zhu, X. 2018. Reinforcement learning for relation classification from noisy data. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- He, Z.; Chen, W.; Li, Z.; Zhang, M.; Zhang, W.; and Zhang, M. 2018. See: Syntax-aware entity embedding for neural relation extraction. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Hoffmann, R.; Zhang, C.; Ling, X.; Zettlemoyer, L.; and Weld, D. S. 2011. Knowledge-based weak supervision for information extraction of overlapping relations. In *Proceedings of ACL*, 541–550.
- Hu, Z.; Ma, X.; Liu, Z.; Hovy, E.; and Xing, E. 2016. Harnessing deep neural networks with logic rules. *arXiv preprint arXiv:1603.06318*.
- Ji, G.; Liu, K.; He, S.; and Zhao, J. 2017. Distant supervision for relation extraction with sentence-level attention and entity descriptions. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Jia, W.; Dai, D.; Xiao, X.; and Wu, H. 2019. Arnor: Attention regularization based noise reduction for distant supervision relation classification. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, 1399–1408.
- Lai, Y.; Feng, Y.; Yu, X.; Wang, Z.; Xu, K.; and Zhao, D. 2019. Lattice cnns for matching based chinese question answering. *arXiv preprint arXiv:1902.09087*.
- Lin, Y.; Shen, S.; Liu, Z.; Luan, H.; and Sun, M. 2016. Neural relation extraction with selective attention over instances. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, 2124–2133.
- Luo, B.; Feng, Y.; Wang, Z.; Huang, S.; Yan, R.; and Zhao, D. 2018. Marrying up regular expressions with neural networks: A case study for spoken language understanding. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, 2083–2093.
- Mintz, M.; Bills, S.; Snow, R.; and Jurafsky, D. 2009. Distant supervision for relation extraction without labeled data. 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: Volume 2-Volume 2*, 1003–1011. Association for Computational Linguistics.
- Qin, P.; Xu, W.; and Wang, W. Y. 2018. Dsgan: generative adversarial training for distant supervision relation extraction. *arXiv preprint arXiv:1805.09929*.
- Riedel, S.; Yao, L.; and McCallum, A. 2010. Modeling relations and their mentions without labeled text. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 148–163. Springer.
- Suchanek, F.; Fan, J.; Hoffmann, R.; Riedel, S.; and Talukdar, P. P. 2013. Advances in automated knowledge base construction. *SIGMOD Records journal, March*.
- Surdeanu, M.; Tibshirani, J.; Nallapati, R.; and Manning, C. D. 2012. Multi-instance multi-label learning for relation extraction. In *Proceedings of the 2012 joint conference on empirical methods in natural language processing and computational natural language learning*, 455–465. Association for Computational Linguistics.
- Vashishth, S.; Joshi, R.; Prayaga, S. S.; Bhattacharyya, C.; and Talukdar, P. 2018. Reside: Improving distantly-supervised neural relation extraction using side information. *arXiv preprint arXiv:1812.04361*.
- Wu, G.; He, Y.; and Hu, X. 2018. Entity linking: an issue to extract corresponding entity with knowledge base. *IEEE Access* 6:6220–6231.
- Xu, J.; Zhang, Z.; Friedman, T.; Liang, Y.; and den Broeck, G. V. 2018. A semantic loss function for deep learning with symbolic knowledge. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, 5498–5507.
- Yu, M.; Yin, W.; Hasan, K. S.; dos Santos, C.; Xiang, B.; and Zhou, B. 2017. Improved neural relation detection for knowledge base question answering. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 571–581. Vancouver, Canada: Association for Computational Linguistics.
- Zeng, D.; Liu, K.; Lai, S.; Zhou, G.; and Zhao, J. 2014. Relation classification via convolutional deep neural network. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, 2335–2344.
- Zeng, D.; Liu, K.; Chen, Y.; and Zhao, J. 2015. Distant supervision for relation extraction via piecewise convolutional neural networks. In *EMNLP*, 1753–1762.