

Integrating Deep Learning with Logic Fusion for Information Extraction

Wenya Wang and Sinno Jialin Pan

Nanyang Technological University, Singapore
{wangwy, sinnopan}@ntu.edu.sg

Abstract

Information extraction (IE) aims to produce structured information from an input text, e.g., *Named Entity Recognition* and *Relation Extraction*. Various attempts have been proposed for IE via feature engineering or deep learning. However, most of them fail to associate the complex relationships inherent in the task itself, which has proven to be especially crucial. For example, the relation between 2 entities is highly dependent on their entity types. These dependencies can be regarded as complex constraints that can be efficiently expressed as logical rules. To combine such logic reasoning capabilities with learning capabilities of deep neural networks, we propose to integrate logical knowledge in the form of first-order logic into a deep learning system, which can be trained jointly in an end-to-end manner. The integrated framework is able to enhance neural outputs with knowledge regularization via logic rules, and at the same time update the weights of logic rules to comply with the characteristics of the training data. We demonstrate the effectiveness and generalization of the proposed model on multiple IE tasks.

Introduction

Information extraction (IE) involves the identification of important information from a piece of input text and is a fundamental step towards knowledge inference. Various problems can be categorized as IE tasks, e.g., *Named Entity Recognition (NER)*, *Entity Linking*, *Opinion Target Extraction (OTE)*, *Relation Extraction (RE)*, etc. In this work, we target at 2 challenging IE tasks including OTE and end-to-end RE. Given an input text, end-to-end RE aims to extract target entities as well as entity relations (Li and Ji 2014). For example, given the sentence “*Rome is in Lazio province and Naples in Campania*”, the task requires the identification of *Rome*, *Lazio*, *Naples* and *Campania* as *location* entities, and the relation between *Rome* and *Lazio* as *Located_In*, same for the relation between *Naples* and *Campania*. The task of OTE aims to identify opinion targets within an opinionated text (Hu and Liu 2004), e.g., *service staff* in “*The service staff in this restaurant is very kind*”.

Deep neural networks (DNNs) have been widely used for various IE tasks. Existing works adopted convolutional

neural networks (Xu et al. 2018a; Adel and Schütze 2017) and recurrent/recursive neural networks (Wang et al. 2016; Miwa and Sasaki 2014) to learn context-aware and high-level features to facilitate predictions. Pointer networks have also been proposed for relation extraction (Katiyar and Cardie 2017). Despite their advantage over low-level feature engineering, the complex networks make learning harder when the amount of training data is insufficient, which is the case for many IE tasks. Moreover, the automation in DNNs makes it challenging to inject prior knowledge to guide the training process. On the opposite, symbolic logic systems provide an effective way to express complex domain knowledge in terms of logic rules and have proven to be advantageous when data is scarce. Inspired by the cognitive process that learns from both experiences and background knowledge, recent years have witnessed a growing interest in combining deep learning with logic reasoning (Manhaeve et al. 2018; Dong et al. 2019) mostly for solving logical problems.

To enhance the extraction performance in the NLP domain, we propose to incorporate domain knowledge as logic rules that are integrated into the representation learning system through a unified framework. The proposed model consists of a deep learning module as well as a logic module, where the deep learning module contains a transformer-style neural network to learn a rich feature representation for each word. The transformer model computes complex word-level correlations in multiple dimensions regardless of context distance (Vaswani et al. 2017), and has shown promising results in several NLP tasks, e.g., semantic role labeling (Tan et al. 2018). We believe this mechanism could be more beneficial to propagate information between related entities, compared to other deep models. The multi-head attention weight indicate the interactions between each pair of words which can be further fed into a relation classifier. The logic module is composed of a set of logic rules represented by First-order Logic (FOL). These rules explicitly specify the complex relationships in the output label space, which could not be handled using simple constraints. For example, a FOL rule, $Live_In(Z, X) \wedge person(Z) \Rightarrow location(X)$, specifies that if the relation between two entities is *Live_In* and the first entity is of type *person*, then the second entity should have type *location*.

To associate distributed features with logic reasoning, we integrate the deep learning module and the logic module through 2 operations: 1) We design some mapping functions such that the information from the neurons could be passed to the logic system. Specifically, the neural outputs are treated as the inputs to the logic module, which combined with probabilistic logic operators, produces the logic outputs. Hence the outputs from the logic module reflects both neural learning and logic interactions among correlated atoms. Furthermore, a learnable weight is assigned to each logic rule to indicate its confidence level. The learnable weight for each rule makes the logic system more flexible and adaptable to specific training dataset, where a higher weight makes the corresponding rule more important within the corpus. 2) A discrepancy loss is proposed to measure the disagreement between the deep learning module and the logic system, which is minimized to allow for regularization of DNNs via logical knowledge. The discrepancy loss prompts the update of neural parameters towards rule-constrained directions, and at the same time adjusts the rule weights to be compatible with specific corpus.

To summarize, the proposed framework has the following contributions: 1) We use transformer mechanism for IE tasks to fully exploit interactions among the **input space**, which is also indicative for relation predictions. 2) We use logic rules to enforce complicated correlations in the **output space** and integrate these rules into the distributed representation learning system with a joint learning mechanism to achieve joint inference. To the best of our knowledge, this is the first work for information extraction that unifies DNN with logic knowledge in a rather *smooth* way to benefit learning of each other. 3) We introduce a general framework for knowledge fusion through discrepancy minimization, which can be adopted in various DNN models. We also demonstrate its effectiveness on different IE tasks.

Related Work

Information Extraction Various approaches have been proposed for entity and relation extraction, either through a pipeline procedure, or a joint inference framework. The pipeline strategy first learns an entity extraction model and then independently predicts relations based on the extracted entities (Chan and Roth 2011; Lin et al. 2016). This strategy suffers from error propagation. To solve this problem, joint inference is proposed to learn shared information between two subtasks by sharing parameters (Miwa and Bansal 2016; Katiyar and Cardie 2017; Bekoulis et al. 2018b; 2018a) and adopting novel tagging scheme to further model task interactions (Roth and Yih 2004; Li and Ji 2014; Miwa and Sasaki 2014; Gupta, Schütze, and Andrassy 2016; Zheng et al. 2017; Zhang, Zhang, and Fu 2017; Adel and Schütze 2017; Sun et al. 2018; Wang et al. 2018; Dai et al. 2019).

For opinion target extraction, existing works either used pre-defined rules/features to identify the targets (Hu and Liu 2004; Qiu et al. 2011; Li et al. 2010), or applied deep learning models with sequence labeling strategy (Yin et al. 2016; Wang et al. 2016; 2017; Li and Lam 2017; Xu et al. 2018a). However, DNNs only implicitly exploit the input interactions, without controlling what is learned. Yu, Jiang, and

Xia (2019) used integer linear programming with explicit constraints for joint inference as a post-processing step.

Deep Learning with Logic Rules The combination of neural learning systems with symbolic rules has long since been proposed, known as neural-symbolic systems (Garcez, Broda, and Gabbay 2012; Manhaeve et al. 2018; Dong et al. 2019; Sourek et al. 2018) that construct a network or connect the distributed systems with given rules for reasoning and inference in logic domains. Xu et al. treated logic knowledge as semantic regularization in the loss function. The injection of logic rules in NLP tasks was recently proposed in (Rocktäschel, Singh, and Riedel 2015; Guo et al. 2016) for relation and knowledge graph learning that embed logic into the same space as distributed features in a single system. Hu et al. (2016) fused logical knowledge into deep models through posterior regularization. Logic rules were also used as evidences to construct adversarial sets (Minervini et al. 2017; Minervini and Riedel 2018), or as a form of indirect supervision (Wang and Poon 2018) to improve model training. Li and Srikumar (2019) augmented deep learning models with logic neurons. In this work, we propose to combine DNN with logic in a smooth way, which adopts probabilistic logic instead of 0/1 hard assignments to facilitate backpropagation through the whole framework. Instead of inserting logic within the DNN architecture, we use a discrepancy loss to progressively bridge the gap between deep learning and logic rules.

Problem Definition & Motivation

We use end-to-end RE, which aims to jointly extract entities and their relations, as a motivating task to describe our proposed method. Denote by \mathcal{E} and \mathcal{R} the set of possible entity types and relation categories, respectively.¹ Given an input sentence $\{w_1, w_2, \dots, w_m\}$, entity extraction involves both entity segmentation as well as entity typing. We use BIO encoding scheme combined with entity types to form the sequence of output labels $y = \{y_1, y_2, \dots, y_m\}$, where $y_i \in \{\text{B-}E_j, \text{I-}E_j, \text{O}\}_{E_j \in \mathcal{E}}$. For example, B-PER (I-PER) indicates the beginning (inside) position of an entity of type *person*. Relation extraction aims to output a set of triplets (e_1, e_2, r) , where e_1 and e_2 represents the first and second entity, respectively, and $r \in \mathcal{R}$ indicates the relation type between them. In this work, we treat entity extraction as a sequence labeling problem and relation extraction as a classification problem based on the identified entities.

A key motivation behind our proposed method is the complex correlation inherent both in the input and output spaces. In the input space, there exist intensive interactions among entities within a sentence to facilitate information propagation. For example, if *Rome* is extracted as a *location* entity and has close relationship with *Lazio*, it may help to identify *Lazio* as another entity. To exploit these interactions, we use the transformer mechanism with multi-head self-attentions to generate a correlation factor for each pair of words, which is injected for both entity and relation predictions. However, DNNs can only implicitly capture some correlations with-

¹Note that the task of OTE can be regarded as a special case of RE, where there is only one entity type.

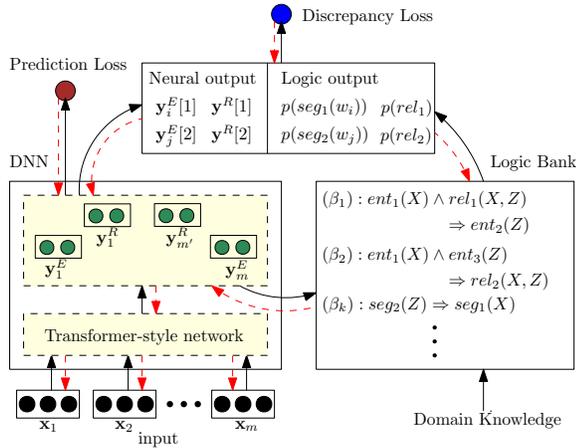


Figure 1: The overall architecture.

out actually enforcing specific relationships. For example, if we know that one of the entities is of type *person* and the relations between the entities is *Live_In*, the other entity should be of type *location*. Even within entity predictions, there also exists dependencies enforcing possible positions within an entity, e.g., label “O” cannot be followed by “I”. These relationships are especially crucial for the task, and yet few existing works have taken them into account. Although sequence labeling models, like Conditional Random Fields (CRFs) (Lafferty, McCallum, and Pereira 2001), are able to encode certain segmentation rules implicitly, they may not learn the optimal strategy when there is insufficient training data. Furthermore, they fail to model more complex relational rules. A few works treat the complex relationships among objects as constraints within the objective function. However, it is non-trivial to express those complex constraints, and the resultant optimization is challenging. Moreover, the constraints cannot be updated to align with the training corpus. In the literature, symbolic logic rules are well-known to be effective at modeling complex semantic relationships. For example, a dependency between entity types and relations can be expressed using FOL as $person(X) \wedge Live_In(X, Z) \Rightarrow location(Z)$. When the training data is insufficient, as is the case for IE, logic rules provide crucial clues to assist learning.

Compared with the pipeline approach which is prone to error propagation, joint inference on related tasks, e.g., entity recognition and relation classification, has shown to be effective for (end-to-end) information extraction. Though joint learning has been proposed in deep architectures, most existing models fail to explicitly enforce the consistency among separate tasks. To address this problem, it is desirable to integrate logic rules specifying task relationships for joint inference. At this point, we propose to unify deep learning with logic rules in an end-to-end learning framework. To align discrete symbolic system with distributed representation learning, we propose to compute logic rules in a probabilistic way and define mapping functions to map the continuous output from a DNN to the logic units. Furthermore, a discrepancy loss is proposed that measure the dis-

crepancy between the DNN outputs and the logic outputs to make these two modules consistent with each other. The discrepancy loss is able to regularize the DNN through domain knowledge, and at the same time update the logic module to comply with the training data.

Methodology

The overall architecture of the proposed method is shown in Figure 1. It consists of 3 components, namely a deep neural network, a logic bank and a discrepancy unit. The DNN component takes a sequence of words as the input and finally produces a prediction vector for each word (and possibly candidate relations). The logic bank is fed with general domain knowledge that is easy to obtain and formalizes the knowledge as a set of first-order logic rules. Note that we assign a non-negative weight to each logic rule to indicate its confidence level which is updated according to the training corpus. The output from DNN is fed into the logic module to produce logic output. The discrepancy unit is responsible for aligning neural outputs with outputs from the logic bank. Specifically, we compute a distance between the output distributions from DNN and logic module with the aim to minimize the distance throughout the learning process.

First-Order Logic

FOL is formed from *constants*, *variables* and *predicates* with propositional connectives \wedge , \vee , \neg and quantifiers. To avoid confusion, we use upper-case letters (A) to represent variables and lower-case letters (a) to represent constants. An *atom* is an n -ary predicate with n arguments ($R(A, B)$). A ground atom assigns constants to all of its arguments ($R(a, b)$). A clause can be written in the form of a rule: $B_1 \wedge \dots \wedge B_k \Rightarrow H$, where H is called the consequent of the rule and $B_1 \wedge \dots \wedge B_k$ is the precondition. The *grounding* of a clause is a substitution that maps each occurring variable in the clause to a constant: $B_{1(\phi)} \wedge \dots \wedge B_{n(\phi)} \Rightarrow H_{(\phi)}$, where ϕ denotes a substitution. A *Herbrand interpretation* is a mapping that assigns a truth value to each ground atom. To make it a *Herbrand model*, all the logic formulas should be satisfied. To find a Herbrand model, a feasible method is through the *immediate consequence operator*, which is a mapping T_p from Herbrand interpretation to itself:

$$T_p(I) = \{H_{(\phi)} \mid (B_1 \wedge \dots \wedge B_n \Rightarrow H) \in \mathcal{P}, \{B_{1(\phi)}, \dots, B_{n(\phi)}\} \in I\}, \quad (1)$$

where I is a Herbrand interpretation, \mathcal{P} is a set of clauses. Given known grounded atoms, we can find other grounded atoms as immediate consequences of the logic formulas. In our formulation, we use neural networks to simulate the immediate consequence operator and applies probabilistic logic where each formula is assigned a confidence score and each grounded atom has a continuous truth value within $[0, 1]$ to indicate its probability of being true.

Deep Learning Module

The deep learning component is modeled as a transformer-style network consisting of multiple layers of self-attentions and Bi-GRU in order to model both sequential and distant

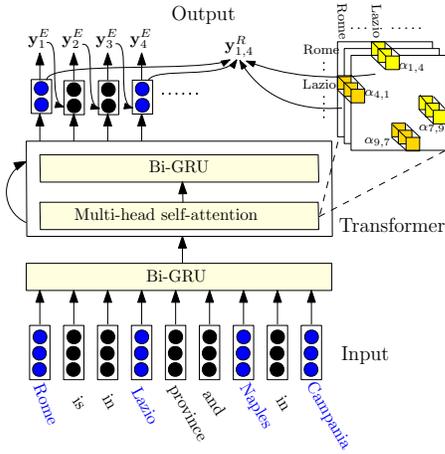


Figure 2: Deep learning module with transformer.

dependencies. A concrete structure is shown in Figure 2. Specifically, given a sequence of input words, we first construct an input embedding $\mathbf{x}_i = [\mathbf{x}_i^e; \mathbf{x}_i^p]$ for each word by concatenating its pre-trained word embedding \mathbf{x}_i^e and its associated POS-tag embedding \mathbf{x}_i^p . To incorporate sequential context interactions, a Bi-GRU model with parameters Θ is firstly applied on top of \mathbf{x}_i to generate hidden representations \mathbf{h}_i by considering both forward and backward information. Mathematically, we denote this process by

$$\mathbf{h}_i = [\vec{\mathbf{h}}_i; \overleftarrow{\mathbf{h}}_i] = [f(\mathbf{x}_i, \vec{\mathbf{h}}_{i-1}; \Theta); f(\mathbf{x}_i, \overleftarrow{\mathbf{h}}_{i+1}; \Theta)],$$

where $\vec{\mathbf{h}}_i$ and $\overleftarrow{\mathbf{h}}_i$ indicate the forward and backward GRU output, respectively.

Subsequently, \mathbf{h}_i is further fed into a transformer model consisting of multiple layers, where each layer stacks a Bi-GRU network on top of a multi-head self-attention module. Specifically, the self-attention network transforms its input \mathbf{h}_i^t at t -th layer to $\tilde{\mathbf{h}}_i^t$ via a series of attention mechanism. For ease of illustration, we drop the superscript t in the sequel. Mathematically, given a query vector corresponding to an input vector \mathbf{h}_i , each attention head computes one type of interactions between itself and other positions within the sequence and produces a transformed hidden representation:

$$\tilde{\mathbf{h}}_i^c = \sum_{j=1}^m \alpha_{ij}^c (\mathbf{W}_v^c \mathbf{h}_j), \text{ and} \quad (2)$$

$$\alpha_i^c = \text{softmax}\left(\frac{(\mathbf{W}_q^c \mathbf{h}_i)(\mathbf{W}_k^c \mathbf{H})}{\sqrt{d}}\right), \quad (3)$$

where \mathbf{H} is a matrix consisting of \mathbf{h}_i as column vectors, and d is the dimension of \mathbf{h}_i . $\{\mathbf{W}_v^c, \mathbf{W}_q^c, \mathbf{W}_k^c\}$ are transformation matrices corresponding to the c -th attention head. We define C different transformations for multi-head mechanism, where each transformation accounts for one representative interaction space. For IE, each head is regarded as computing a different relation between 2 words. A final hidden vector is produced via $\tilde{\mathbf{h}}_i = \mathbf{W}[\tilde{\mathbf{h}}_i^1; \dots; \tilde{\mathbf{h}}_i^C]$.

Denote the final feature representation after applying Bi-GRU in the last transformer layer T as \mathbf{h}^T . The neural out-

puts for entity prediction \mathbf{y}^E are generated through a fully-connected layer followed by a softmax layer:

$$\mathbf{s}_i^E = \tanh(\mathbf{W}_h^E [\mathbf{h}_i^T; \mathbf{x}_{i-1}^l] + \mathbf{b}_h^E), \quad (4)$$

$$p(\mathbf{y}_i^E | \mathbf{x}) = \text{softmax}(\mathbf{W}_y^E \mathbf{s}_i^E + \mathbf{b}_y^E), \quad (5)$$

where \mathbf{x}_{i-1}^l denotes the entity label embedding of the previous token. The injection of label embeddings implicitly informs entity label dependencies. Similarly, the relation prediction \mathbf{y}^R is produced for each pair of extracted entities via

$$\mathbf{s}_{i,j}^R = \tanh(\mathbf{W}_h^R [\mathbf{v}_i; \mathbf{v}_j; \alpha_{i,j}; \alpha_{j,i}] + \mathbf{b}_h^R), \quad (6)$$

$$p(\mathbf{y}_{i,j}^R | \mathbf{x}, \mathcal{E}) = \text{softmax}(\mathbf{W}_y^R \mathbf{s}_{i,j}^R + \mathbf{b}_y^R), \quad (7)$$

where $\mathbf{v}_i = [\mathbf{h}_i^T; \mathbf{x}_i^l]$, \mathcal{E} indicates the set of extracted entities, and $\alpha_{i,j}$ represents the multi-head attention weight vector between w_i and w_j . The attention vectors contain explicit correlation information that could assist relation prediction, as will be verified later in experiments.

Logic Fusion via Discrepancy Loss

As discussed, the DNN module introduced in the previous section is able to exploit word-level interactions implicitly via feature learning, but fails to consider more complex relationships among outputs. For example, the relationship between entity types and their relations, which are non-trivial to be injected into DNNs, but can be effectively formulated as logical rules. We construct 2 types of FOLs to specify the relationships among entity and relation labels. The first type focuses on the dependencies of segmentation labels, denoted by $\text{seg}_b(Z) \Rightarrow \text{seg}_a(X)$, which means if the segmentation label for variable Z is b , the segmentation label for variable X is deduced as a , with $a, b \in \{\text{B}, \text{I}, \text{O}\}$. An example is $\text{seg}_B(w_i) \Rightarrow \text{seg}_O(w_{i-1})$ that enforces the previous word to have label O when the current word is the beginning of an entity. The second rule models the correlations between entity types and relations, e.g., $\text{entity}_c(X) \wedge \text{rel}_l(X, Z) \Rightarrow \text{entity}_d(Z)$, which means relation l only appears when the first and second entity has type c and d , respectively.

To make logic rules compatible with DNN outputs, we adopt probabilistic logic where each atom is assigned a continuous value in $[0, 1]$ and convert Boolean operations to work on probabilistic units. Hence, we define the following:

Definition. A mapping Γ from the language of FOL \mathcal{L} to the continuous space \mathbb{R} ($\Gamma: \mathcal{L} \rightarrow \mathbb{R}$) is defined as follows,

- $\Gamma(P) = \mathbf{y}(P)$ with P representing a grounded atom and \mathbf{y} as the neural output. For example, $\Gamma(\text{entity}_c(w_i)) = \mathbf{y}_i^E[c]$, where c indicates the index of the entity type. This can be interpreted as: the probability for w_i belonging to entity_c equals to its corresponding neural output $\mathbf{y}_i^E[c]$.
- $\Gamma(P_1 \wedge \dots \wedge P_n) = \sigma(a_0(\sum_{i=1}^n \Gamma(P_i) - n) + b_0)$.
- $\Gamma(P_1 \vee \dots \vee P_n) = \sigma(a_1 \sum_{i=1}^n \Gamma(P_i) + b_1)$.
- $\Gamma(\neg P) = 1 - \Gamma(P)$.

Here σ indicates the sigmoid activation function. The last 3 mappings are able to approximate the semantics of logic operators according to (Sourek et al. 2018). With these mapping functions, we compute a soft-version of the *immediate*

consequence operator using (1) to produce the output $Y_{\mathcal{L}}$ from the logic module. Specifically, the value of the consequent atom $H_{(\phi)}$ in each rule is deduced by applying Γ on the rule body $B_{(\phi)}$, given a grounded clause $B_{(\phi)} \Rightarrow H_{(\phi)}$, where $B_{(\phi)}$ denotes the conjunction $B_{1(\phi)} \wedge \dots \wedge B_{n(\phi)}$.

$$\begin{aligned} Y_{\mathcal{L}}(H_{(\phi)}) &= \Gamma(B_{1(\phi)} \wedge \dots \wedge B_{n(\phi)}) \\ &= \sigma(a_0(\sum_{i=1}^n \Gamma(B_{i(\phi)}) - n) + b_0) \\ &= \sigma(a_0(\sum_{i=1}^n \mathbf{y}(B_{i(\phi)}) - n) + b_0). \end{aligned} \quad (8)$$

A detailed procedure to produce logic output is shown in Algorithm 1. Given neural outputs $\{\mathbf{y}_i^E\}_{i=1}^m$ in (5) and $\{\mathbf{y}_l^R\}_{l=1}^{m'}$ in (7) for each sentence, and a set of rules $\{B^{(k)} \Rightarrow H^{(k)}\}_{k=1}^K$, the logic system produces a logic output $u_{i,k}^E$ for each word and $u_{l,k}^R$ for each relation corresponding to each rule r_k as following: For each rule $B^{(k)} \Rightarrow H^{(k)}$, we find its satisfying groundings ϕ for $B^{(k)}$ from neural predictions and generate the logic output $Y_{\mathcal{L}}(H_{(\phi)})$ for its consequent atom $H_{(\phi)}^{(k)}$ using (8). For example, given the rule $entity_c(X) \wedge rel_l(X, Z) \Rightarrow entity_d(Z)$, if the neural model predicts w_i as $entity_c$ and the relation between w_i and w_j as rel_l , ($X = w_i, Z = w_j$) is regarded as a satisfying grounding. Then the logic output for the consequent atom $entity_d(w_j)$ will be produced as $\sigma(a_0(\mathbf{y}_i^E[c] + \mathbf{y}_{i,j}^R[l] - 2) + b_0)$. We use Φ^k and Γ^k to collect the deduced consequent groundings and their logic values, respectively. The final logic output for each word $u_{i,k}^E$ (relation $u_{l,k}^R$) for each rule r_k is obtained by aggregating the logic values across all the situations when acting as a consequent atom.

By making the DNN and the logic module compatible with each other, we can measure their discrepancy $\ell_D(\mathcal{F}, \mathcal{L})$ by comparing the distributions of their outputs:

$$\begin{aligned} \ell_D(\mathcal{F}, \mathcal{L}) &= \mathbb{E}_{x \sim \mathcal{X}}(d(Y_{\mathcal{F}}(x), Y_{\mathcal{L}}(x))) \\ &= \frac{1}{K} \sum_{\{B^{(k)} \Rightarrow H^{(k)}\}} \frac{1}{|\Phi^k|} \sum_{\phi \in \Phi^k} \beta_k d(\mathbf{y}(\phi), \mathbf{u}_k(\phi)), \end{aligned} \quad (9)$$

where $Y_{\mathcal{F}}(x)$ and $Y_{\mathcal{L}}(x)$ denote the neural output and logic output, respectively. Φ^k collects the consequent atoms whose precondition is satisfied. We further assign a confidence weight $\beta_k \in [0, 1]$ for each rule to indicate its confidence and adjust its contribution to the discrepancy loss. The higher the weight, the more penalty to be given when neural outputs disagree with logic outputs. We use Mean-Squared-Error as the distance metric $d(\cdot, \cdot)$, because it provides a better gradient flow for a more stable training process.

Training

The integrated model can be trained end-to-end with gradient descent by minimizing $\ell = \ell_Y + \ell_D$, where ℓ_Y is the prediction loss for the deep learning model. Here we use cross-entropy loss for both entity and relation predictions:

$$\ell_Y = -\frac{1}{N} \sum_{n=1}^N (\log p(\hat{\mathbf{y}}_n^E | \mathbf{x}_n) + \log p(\hat{\mathbf{y}}_n^R | \mathbf{x}_n, \mathcal{E}_n)), \quad (10)$$

Algorithm 1 Deep Logic

Input: Neural softmax outputs $\{\mathbf{y}_i^E\}_{i=1}^m$ (entities) and $\{\mathbf{y}_l^R\}_{l=1}^{m'}$ (relations) for each sentence.
Output: $\{u_{i,k}^E\}_{i=1,k=1}^{m,K}$ (entities) and $\{u_{l,k}^R\}_{l=1,k=1}^{m',K}$ (relations)
Initialize: $u_{i,k}^E = 0, u_{l,k}^R = 0$ for $i \in \{1, \dots, m\}, l \in \{1, \dots, m'\}, k \in \{1, \dots, K\}$.
 Collect feasible groundings of rule head and rule body.
for each rule $r_k : B^{(k)} \Rightarrow H^{(k)}$ **do**
 1: Initialize $\Phi^k = \{\}, \Gamma^k = \{\}$
 2: Find a grounding ϕ such that each $B_{j(\phi)}^{(k)}$ in $B_{(\phi)}^{(k)}$ is true according to neural predictions $\{\mathbf{y}_i^E\}_{i=1}^m, \{\mathbf{y}_l^R\}_{l=1}^{m'}$
 3: Update $\Phi^k \leftarrow \Phi^k \cup \{H_{(\phi)}^{(k)}\}, \Gamma^k \leftarrow \Gamma^k \cup \{Y_{\mathcal{L}}(H_{(\phi)}^{(k)})\}$
end for
 Compute logic output for rule heads
for k from 1 to K **do**
 Initialize $c_i^E = 0$ for $i \in \{1, \dots, m\}, c_l^R = 0$ for $l \in \{1, \dots, m'\}$
 for $(\phi, \gamma) \in (\Phi^k, \Gamma^k)$ **do**
 Return the exact word w_i or relation r_l that corresponds to grounding ϕ
 Update $u_{i,k}^E \leftarrow u_{i,k}^E + \gamma, c_i^E \leftarrow c_i^E + 1$ or $u_{l,k}^R \leftarrow u_{l,k}^R + \gamma, c_l^R \leftarrow c_l^R + 1$
 end for
 $u_{i,k}^E \leftarrow u_{i,k}^E / c_i^E, u_{l,k}^R \leftarrow u_{l,k}^R / c_l^R$
end for

where $p(\hat{\mathbf{y}}_n^E | \mathbf{x}_n) = \prod_{i=1}^{|\mathcal{S}_n|} p(\mathbf{y}_i^E = \hat{\mathbf{y}}_i^E | \mathbf{x}_n)$ using (5). $|\mathcal{S}_n|$ indicates the length of the n th sentence. Similar procedure applies to $p(\hat{\mathbf{y}}_n^R | \mathbf{x}_n, \mathcal{E}_n)$ using (7). \mathcal{E}_n denotes the set of extracted entities for the n -th sequence. The backpropagation procedure is revealed in Figure 1 via (dashed) downward arrows. Specifically, the discrepancy loss updates both neural and logic outputs, together with rule weights according to (9) through gradient descent. To restrict the logic weights within $[0, 1]$, we apply a sigmoid function such that $\beta_k = \sigma(\beta'_k)$. Then the gradient of the logic weight becomes

$$\frac{\partial \ell_D}{\partial \beta'_k} = \frac{1}{K |\Phi^k|} \sum_{\phi \in \Phi^k} d\sigma(\beta'_k) (1 - \sigma(\beta'_k)), \quad (11)$$

where $d = d(\mathbf{y}(\phi), \mathbf{u}_k(\phi))$. The gradients for logic output $\mathbf{u}_k(\phi)$ is further passed back to neural logits $\mathbf{y}(B_{(\phi)})$ according to (8), which combined with the classification loss, updates all the parameters within the neural model. Ideally, the discrepancy loss will punish the situation when the neural output highly differs from the logic output. In this case, the deep model will modify its network to be more aligned with the logic rules. On the other hand, the logic module will adapt its weights as well as the mappings that are passed back to the neurons. For example, if the deep module predicts *Rome* as *location*, *Lazio* as *location* and their relation as *OrgBased_In* (which is wrong). When feeding them into the rule $loc(Rome) \wedge OrgBased_In(Lazio, Rome) \Rightarrow org(Lazio)$, the logic output for $entity_{org}(Lazio)$ will be high, different from the neural output. In this case, the discrepancy revises its rule body to decrease the incorrect neural output for relation *OrgBased_In*.

	Restaurant14	Laptop14	Restaurant16
(Wang et al. 2016)	84.25	77.26	69.74
(Wang et al. 2017)	84.38	76.45	73.87
(Li and Lam 2017)	-	77.58	73.44
(Xu et al. 2018a)	84.24	81.59	74.37
(Yu, Jiang, and Xia 2019)	84.50	78.69	-
TransF	84.64	81.76	73.56
Ours	85.62	82.46	74.67

Table 1: Comparison with baselines on OTE.

Experiment

To demonstrate the effectiveness of our proposed method, we conduct experiments on 5 datasets from 2 tasks:

OTE: We use Restaurant and Laptop reviews from SemEval 2014 and 2016 (Pontiki et al. 2014; 2016).

End-to-End RE: 1) **TREC**: entity and relation dataset introduced in (Roth and Yih 2004). It has 4 entity types: *others*, *person*, *location* and *organization*, and 5 relations: *Located_In*, *Live_In*, *OrgBased_In*, *Work_For* and *Kill*. We follow the preprocessing from (Gupta, Schütze, and Andrassy 2016) 2) **ACE05**: annotated data with 7 coarse-grained entity types and 6 coarse-grained relation types between entities. We follow the same setting as (Li and Ji 2014).

For the OTE task, we follow the setting in (Wang et al. 2016) by first pre-training the word embedding using *word2vec* (Mikolov et al. 2013) on Yelp Challenge dataset² and electronic dataset in Amazon reviews³ for restaurant domain and laptop domain, respectively. For RE task, the word embedding is pre-trained on wikipedia corpus using *Glove* (Pennington, Socher, and Manning 2014). For all experiments, the dimension for word embedding and POS embedding is set to 300 and 50, respectively. The hidden layers has dimension 200. We set label embedding with dimension 25. Following (Vaswani et al. 2017), we also use positional encoding that is added to the input vectors. The multi-head self-attentions adopts 10 heads that leads to 10-dim attention weight vectors. For RE task, we use scheduled sampling, similar to (Miwa and Bansal 2016). To train the model, adadelta is adopted with initial rate as 1.0 and with dropout rate 0.1. For evaluation, we use micro-F1 scores on non-negative classes. An entity is counted as correct based on exact match. A relation is correct if both of its entities are correct and the relation type matches the ground-truth label.

Results & Analysis

Comparison on OTE task: Table 1 shows the comparison results for opinion target extraction with popular baselines. The last two rows indicate our proposed models, where *TransF* is the deep learning module without logic integration. Since the OTE task can be viewed as single-class entity extraction, the proposed model can be adapted to this task by ignoring relation predictions. From the results, we can observe that even without logic rules, the transformer model is able to achieve 2-out-of-3 best results compared to existing works. This proves the effectiveness of transformer for implicit interaction modeling. When considering

²http://www.yelp.com/dataset_challenge

³<http://jmcauley.ucsd.edu/data/amazon/links.html>

Setting	Model	Evaluation	Entity	Relation
w/ Boundary	(Gupta et al. 2016)	relaxed	92.4	69.9
	(Bekoulis et al. 2018b)	relaxed	93.3	67.0
	(Bekoulis et al. 2018a)	relaxed	93.0	68.0
	(Miwa and Sasaki 2014)	strict	92.3	71.0
	(Adel and Schütze 2017)	strict	92.1	65.3
	Pipeline	strict	94.1	69.7
	Pipeline+feat	strict	94.5	70.2
	TransF	strict	94.6	72.8
	Ours (w/o) POS	strict	94.3	71.7
	Ours	strict	95.1	74.1
w/o Boundary	(Miwa and Sasaki 2014)	relaxed	80.7	61.0
	(Adel and Schütze 2017)	relaxed	82.1	62.5
	(Bekoulis et al. 2018b)	strict	83.0	61.0
	(Bekoulis et al. 2018a)	strict	83.6	62.0
	Pipeline	strict	84.2	57.2
	Pipeline+feat	strict	84.2	58.4
	TransF	strict	85.8	62.7
	Ours (w/o) POS	strict	85.0	63.1
	Ours	strict	87.1	64.6

Table 2: Results of End-to-End RE on TREC.

Settings	Models	Res14	Res16	Lap14	TREC		ACE05	
		OTE			Entity	Relation	Entity	Relation
Entity	TransF- α^1	84.1	72.7	77.6	83.3	-	82.9	-
	TransF	84.6	73.6	81.8	84.2	-	83.2	-
	TransF+SR	85.0	73.9	82.1	84.9	-	83.2	-
Joint	TransF- α	-	-	-	84.3	60.6	83.3	57.4
	TransF	-	-	-	85.8	62.7	83.4	59.1
	TransF+SR+RR	85.6	74.7	82.5	87.1	64.6	83.4	59.3

Table 4: Comparisons on different model settings.

logic knowledge, we use segmentation rules that enforces possible segmentation labels for 2 adjacent words. Furthermore, we also incorporate implicit relational rules that state $entity(X) \wedge entity(Z) \Rightarrow related(X, Z)$. This is achieved by keeping the relation prediction layer in the deep learning module. We find this strategy slightly improves our results which will be shown later. The results also show that the integration of logic rules is more effective than CRF by comparing with (Wang et al. 2016). Although Yu, Jiang, and Xia (2019) incorporated explicit constraints through integer linear programming, the separation from DNN during learning makes it suboptimal. This demonstrates the advantage of our unified framework that associates logic reasoning with representation learning. Clearly, our model achieves the state-of-the-art results on all 3 datasets.

Comparison on RE task: The comparison results on TREC is shown in Table 2. Existing works for this domain consists of 2 different settings and evaluations. The first setting is introduced in (Roth and Yih 2004) that assumes the entity boundaries are given and the task is to predict the entity types and relations. The second setting requires both segmentation and entity typing. Evaluations include relaxed version where an entity is regarded as correct if at least one of its consisting words have the correct type prediction. The strict version only treats a predicted entity as correct given a complete match. When boundaries are given, our model could be easily modified to treat each entity as a single unit for type predictions. To show the effect of joint inference with logic rules, we construct a pipeline model by first predicting entities followed by relation predictions given fixed

	(Li and Ji 14)	(Miwa and Bansal 16)	(Zhang et al. 17)	(Sun et al. 18)	(Sun et al. 18) _G	Pipeline	Pipeline+feat	TransF	TransF _G	Ours	Ours _G
Entity	80.8	83.4	83.5	83.4	83.6	83.2	83.2	83.4	83.5	83.4	83.6
Relation	49.5	55.6	57.5	57.8	59.6	57.4	55.3	59.1	59.1	59.3	59.4

Table 3: Comparison with baselines on relation extraction using ACE2005 dataset.

entity parameters. Another model (Pipeline+feat) further appends rule-based features as 1-hot vectors for relation prediction. Obviously, the pipeline model achieves inferior performance and simple features are far less expressive than logic rules. For both settings, our model achieves the best results with a large margin. Furthermore, we also test our model without the POS embedding, shown as ‘‘Ours (w/o POS)’’, which still demonstrates some performance gain.

Table 3 shows the comparison results on ACE05 dataset. Note that Sun et al. (2018) used a non-decomposable global loss with normal local loss to jointly train entity and relation extraction. We denote by (Sun et al. 18) and (Sun et al. 18)_G the model without and with the global loss, respectively. We also implemented the global loss as a fine-tuning step in our method, denoted by the subscript *G*. It is observed that *TransF* achieves the best result for relation extraction and comparable result for entity extraction without sophisticated global training strategy. Global loss does not provide clear benefit for our model. This might indicate that our model already explores global interactions. (Sun et al. 18)_G requires substantial pretraining processes (larger than 500 epochs in our implementation) using the normal loss. As a comparison, our model is more convenient to implement and provides faster convergence, as will be shown later. The logic rules only show slight improvement. We deduce the reason to be the characteristics of the data. Indeed, there are very few absolute relationships between entity types and relations. In most cases, two related entity types can have multiple relation categories, making explicit rules less useful.

Comparison for different model settings: To demonstrate the effect of each component, we conduct experiments on different model settings, with the statistics shown in Table 4. The objective is to analyze the effect of interactions within entities themselves, as well as the correlations between entities and relations. We first remove the relation prediction component to examine the results for solely entity extraction, denoted as ‘Entity’. We construct 3 different settings to compare the results among standard transformer model, removing label embedding (TransF−*x*^l), and integrating only segmentation rules (TransF+SR) with discrepancy loss. As shown, the label embedding slightly improves all the performances which proves its ability to implicitly exploit the segmentation dependencies. However, more improvements are observed by explicitly incorporating logic rules. When jointly predicting entities and relations, we could observe slight improvements upon entity extraction using transformer. This shows positive correlations between entities and their relations. For this joint task, we examine the effect of distributed features (TransF−*α*) as well as symbolic rules (TransF+SR+RR) on the final results. By removing attention weight features, TransF−*α* achieves inferior results on relation predictions, demonstrating the advantage of attentions for relation modeling. More significant

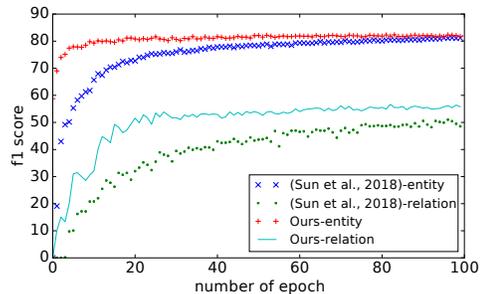


Figure 3: Performance trend comparison on 2 models.

improvement could be observed when incorporating the relationships between entities and relations as explicit rules.

Convergence Comparison: (Sun et al. 18)_G requires a great amount of pre-training epochs using the normal local loss (at least 500 epochs) for the base model to reach reported performance. As a comparison, our model with the normal loss converges much faster, as demonstrated in Figure 3 on the test data. Compared with (Sun et al. 2018), our model reaches to a relatively high performance within 20 epochs. The result for entity extraction has faster convergence rate compared to relation prediction but stays in a rather stable state, which is comparable to (Sun et al. 2018) in subsequent epochs. The relation prediction performance for our model substantially increases in the first place, while still keeps improving later on. We can observe that within 100 epochs, our model achieves clear performance gain over (Sun et al. 2018) on relation extraction.

Conclusion

In this work, we propose a novel unified model to associate distributed learning with symbolic rules. The integrated framework is able to pass information from the neural model to the logic module and compute a discrepancy loss between these two components, which is minimized to update the whole network. The marriage between these two systems could regularize deep learning in the form of knowledge distillation. On the other hand, the logic system is also updated in terms of rule weights to adapt to specific data domain. Experimental results demonstrate the advantage of combining DNNs and logic for joint inference.

Acknowledgements

This work is supported by NTU Nanyang Assistant Professorship (NAP) grant M4081532.020, Singapore MOE AcRF Tier-2 Grant MOE2016-T2-2-060, and the Lee Kuan Yew Postdoctoral Fellowship program, Singapore.

References

- Adel, H., and Schütze, H. 2017. Global normalization of convolutional neural networks for joint entity and relation classification. In *EMNLP*, 1723–1729.
- Bekoulis, I.; Deleu, J.; Demeester, T.; and Develder, C. 2018a. Adversarial training for multi-context joint entity and relation extraction. In *EMNLP*, 1–7.
- Bekoulis, I.; Deleu, J.; Demeester, T.; and Develder, C. 2018b. Joint entity recognition and relation extraction as a multi-head selection problem. *EXPERT SYSTEMS WITH APPLICATIONS* 114:34–45.
- Chan, Y. S., and Roth, D. 2011. Exploiting syntactico-semantic structures for relation extraction. In *ACL-HLT*.
- Dai, D.; Xiao, X.; Lyu, Y.; Dou, S.; She, Q.; and Wang, H. 2019. Joint extraction of entities and overlapping relations using position-attentive sequence labeling. In *AAAI*, 6300–6308.
- Dong, H.; Mao, J.; Lin, T.; Wang, C.; Li, L.; and Zhou, D. 2019. Neural logic machines. In *ICLR*.
- Garcez, A. S. d.; Broda, K. B.; and Gabbay, D. M. 2012. *Neural-Symbolic Learning System: Foundations and Applications*. Springer Science & Business Media.
- Guo, S.; Wang, Q.; Wang, L.; Wang, B.; and Guo, L. 2016. Jointly embedding knowledge graphs and logical rules. In *EMNLP*.
- Gupta, P.; Schütze, H.; and Andrassy, B. 2016. Table filling multi-task recurrent neural network for joint entity and relation extraction. In *COLING*, 2537–2547.
- Hu, M., and Liu, B. 2004. Mining and summarizing customer reviews. In *KDD*, 168–177.
- Hu, Z.; Ma, X.; Liu, Z.; Hovy, E.; and Xing, E. 2016. Harnessing deep neural networks with logic rules. In *ACL*, 2410–2420.
- Katiyar, A., and Cardie, C. 2017. Going out on a limb: Joint extraction of entity mentions and relations without dependency trees. In *ACL*, 917–928.
- Lafferty, J. D.; McCallum, A.; and Pereira, F. C. N. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *ICML*, 282–289.
- Li, Q., and Ji, H. 2014. Incremental joint extraction of entity mentions and relations. In *ACL*, 402–412.
- Li, X., and Lam, W. 2017. Deep multi-task learning for aspect term extraction with memory interaction. In *EMNLP*, 2886–2892.
- Li, T., and Srikumar, V. 2019. Augmenting neural networks with first-order logic. In *ACL*, 292–302.
- Li, F.; Han, C.; Huang, M.; Zhu, X.; Xia, Y.-J.; Zhang, S.; and Yu, H. 2010. Structure-aware review mining and summarization. In *COLING*, 653–661.
- Lin, Y.; Shen, S.; Liu, Z.; Luan, H.; and Sun, M. 2016. Neural relation extraction with selective attention over instances. In *ACL*.
- Manhaeve, R.; Dumancic, S.; Kimmig, A.; Demeester, T.; and De Raedt, L. 2018. Deepproblog: Neural probabilistic logic programming. In *NeurIPS*. 3749–3759.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*, 3111–3119.
- Minervini, P., and Riedel, S. 2018. Adversarially regularising neural NLI models to integrate logical background knowledge. In *CoNLL*, 65–74.
- Minervini, P.; Demeester, T.; Rocktäschel, T.; and Riedel, S. 2017. Adversarial sets for regularised neural link predictors. In *UAI*.
- Miwa, M., and Bansal, M. 2016. End-to-end relation extraction using LSTMs on sequences and tree structures. In *ACL*.
- Miwa, M., and Sasaki, Y. 2014. Modeling joint entity and relation extraction with table representation. In *EMNLP*.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. Glove: Global vectors for word representation. In *EMNLP*, volume 14, 1532–1543.
- Pontiki, M.; Galanis, D.; Pavlopoulos, J.; Papageorgiou, H.; Androutsopoulos, I.; and Manandhar, S. 2014. Semeval-2014 task 4: Aspect based sentiment analysis. In *SemEval*.
- Pontiki, M.; Galanis, D.; Papageorgiou, H.; Androutsopoulos, I.; Manandhar, S.; AL-Smadi, M.; Al-Ayyoub, M.; Zhao, Y.; Qin, B.; Clercq, O. D.; Hoste, V.; Apidianaki, M.; Tannier, X.; Loukachevitch, N.; Kotelnikov, E.; Bel, N.; Jimnez-Zafra, S. M.; and Eryit, G. 2016. In *SemEval 2016*.
- Qiu, G.; Liu, B.; Bu, J.; and Chen, C. 2011. Opinion word expansion and target extraction through double propagation. *Comput. Linguist.* 37(1):9–27.
- Rocktäschel, T.; Singh, S.; and Riedel, S. 2015. Injecting logical background knowledge into embeddings for relation extraction. In *NAACL*, 1119–1129.
- Roth, D., and Yih, W.-t. 2004. A linear programming formulation for global inference in natural language tasks. In *HLT-NAACL 2004 Workshop: CoNLL-2004*, 1–8.
- Sourek, G.; Aschenbrenner, V.; Zelezný, F.; Schockaert, S.; and Kuzelka, O. 2018. Lifted relational neural networks. *JAIR* 62.
- Sun, C.; Wu, Y.; Lan, M.; Sun, S.; Wang, W.; Lee, K.-C.; and Wu, K. 2018. Extracting entities and relations with joint minimum risk training. In *EMNLP*, 2256–2265.
- Tan, Z.; Wang, M.; Xie, J.; Chen, Y.; and Shi, X. 2018. Deep semantic role labeling with self-attention. In *AAAI*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is all you need. In *NIPS*. 5998–6008.
- Wang, H., and Poon, H. 2018. Deep probabilistic logic: A unifying framework for indirect supervision. In *EMNLP*, 1891–1902.
- Wang, W.; Pan, S. J.; Dahlmeier, D.; and Xiao, X. 2016. Recursive neural conditional random fields for aspect-based sentiment analysis. In *EMNLP*.
- Wang, W.; Pan, S. J.; Dahlmeier, D.; and Xiao, X. 2017. Coupled multi-layer tensor network for co-extraction of aspect and opinion terms. In *AAAI*.
- Wang, S.; Zhang, Y.; Che, W.; and Liu, T. 2018. Joint extraction of entities and relations based on a novel graph scheme. In *IJCAI*, 4461–4467.
- Xu, H.; Liu, B.; Shu, L.; and Yu, P. S. 2018a. Double embeddings and CNN-based sequence labeling for aspect extraction. In *ACL*, 592–598.
- Xu, J.; Zhang, Z.; Friedman, T.; Liang, Y.; and Van den Broeck, G. 2018b. A semantic loss function for deep learning with symbolic knowledge. In *ICML*, 5502–5511.
- Yin, Y.; Wei, F.; Dong, L.; Xu, K.; Zhang, M.; and Zhou, M. 2016. Unsupervised word and dependency path embeddings for aspect term extraction. In *IJCAI*.
- Yu, J.; Jiang, J.; and Xia, R. 2019. Global inference for aspect and opinion terms co-extraction based on multi-task neural networks. *IEEE/ACM Trans. Audio, Speech and Lang. Proc.* 27(1):168–177.
- Zhang, M.; Zhang, Y.; and Fu, G. 2017. End-to-end neural relation extraction with global optimization. In *EMNLP*.
- Zheng, S.; Wang, F.; Bao, H.; Hao, Y.; Zhou, P.; and Xu, B. 2017. Joint extraction of entities and relations based on a novel tagging scheme. In *ACL*, 1227–1236.