

Knowledge-guided Open Attribute Value Extraction with Reinforcement Learning

Ye Liu^{1*} Sheng Zhang^{1*} Rui Song^{1†} Suo Feng² Yanghua Xiao^{2†}

¹Department of Statistics, North Carolina State University, Raleigh, NC 27695-8206

²Shanghai Key Laboratory of Data Science, School of Computer Science, Fudan University, China 200433
{yliu87, szhang37, rsong}@ncsu.edu, {fengs17, shawyh}@fudan.edu.cn

Abstract

Open attribute value extraction for emerging entities is an important but challenging task. A lot of previous works formulate the problem as a *question-answering* (QA) task. While the collections of articles from web corpus provide updated information about the emerging entities, the retrieved texts can be noisy, irrelevant, thus leading to inaccurate answers. Effectively filtering out noisy articles as well as bad answers is the key to improving extraction accuracy. Knowledge graph (KG), which contains rich, well organized information about entities, provides a good resource to address the challenge. In this work, we propose a knowledge-guided reinforcement learning (RL) framework for open attribute value extraction. Informed by relevant knowledge in KG, we trained a deep Q-network to sequentially compare extracted answers to improve extraction accuracy. The proposed framework is applicable to different information extraction system. Our experimental results show that our method outperforms the baselines by 16.5 - 27.8%.

1 Introduction

Numerous entities are emerging everyday. The attributes of the entities are often noisy or incomplete, even missing. In the field of electronic commerce, target attributes (e.g., brand, flavor, smell) of new products are often missing (Zheng et al., 2018). In medical analysis, attributes like transmission, genetics and origins of a novel virus are often unknown to people. Even in DBpedia, a well-constructed and large-scale knowledge base extracted from Wikipedia, half of the entities contain less than 5 relationships (Shi and Weninger,

2018). A method that is capable of supplementing reliable attribute values for emerging entities can be highly useful in many applications.

Although information extraction methods have been extensively studied, the task of open attribute value extraction remains challenging. First, the emerging entities may have new attribute values that are absent in the existing KG. Under such circumstances, the prediction methods under the closed-world assumption and the methods that cannot utilize external information are not well suited due to their limited recalls. Second, while web corpus can be used as a good resource to provide relatively updated and relevant articles for large varieties of emerging entities, the articles retrieved from web corpus can be noisy and/or irrelevant, which in turn leads to a limited precision. Finally, even when articles are relevant, the extracted answers might still be inaccurate due to the error-prone information extraction model.

To effectively filter out noisy answers that are obtained either due to the irreverent articles or the errors incurred by the information extraction system, we pose the following two questions: First, how many articles should we collect from the enormous web corpus? Second, how to select the most reliable value out of the pool of all the possible answers extracted from the articles?

There is no common answer to the first question that works for all triplets because of the inconsistent degrees of difficulties in finding the correct attribute values. The decision of when to stop querying more external articles needs to be made after successive evaluations of the candidate answers. Thus the decision making process is inherently sequential.

Reinforcement learning (RL) is a commonly adopted method to deal with sequential decision problems and has been widely studied in the field of robotic and game (Sutton et al., 1998). But

*Ye Liu and Sheng Zhang contributed equally.

†Rui Song and Yanghua Xiao are corresponding authors. Yanghua Xiao was supported by Shanghai Science and Technology Innovation Action Plan (No.19511120400).

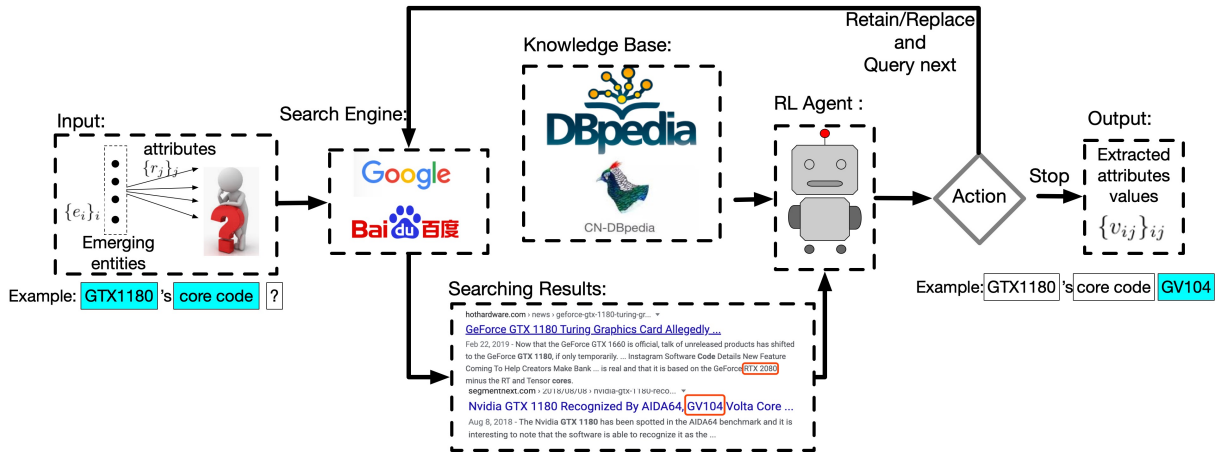


Figure 1: Illustration of overall process. The inputs are pairs of entities and attributes. Relevant articles are retrieved via search engines. The articles together with the KG are fed into the RL agent to inform the selection between candidate answers and the stopping decisions. When the RL agent decides to stop, it will output the best extracted answer.

there are not many researches on open attribute value extraction with RL. One existing literature of RL-based method for value extraction is proposed by (Narasimhan et al., 2016). In their work, a RL framework is designed to improve accuracy of event-related value extraction by acquiring and incorporating external evidences. However, their approach requires a great amount of context information about the specific event of interest during the training process. It is not trivial to extend their framework for open attribute value extraction, because we would need to collect context words and train a new model with annotated data for each emerging attribute. Therefore, their framework cannot be generalized to open attribute value extraction task when various entities and attributes are involved.

While using the context words to construct the states in RL is not suitable in our task, our solution is to leverage the rich, well-organized information in KG, which is not only informative but also generalizable. Such information can be leveraged in answer comparisons, which addresses our second question. For example, to fill the incomplete triplet $\langle iPhone\ 11, display\ resolution, ? \rangle$, from the KG we may find that the attribute values “display resolutions” of an entity that is under category “Phone” is commonly expressed in the format of “xxx by xxxx Pixels”, where x stands for some digit. The typical instances of the attribute values for entities under the same category provide valuable background information.

In this paper, we propose a knowledge-guided

RL framework to perform open attribute value extraction. The RL agent is trained to make good actions for answer selection and stopping time decision. Our experiments show that the proposed framework significantly boosts the extraction performance.

To the best of our knowledge, we are the first to integrate KG in a RL framework to perform open attribute value extraction. In summary, our contribution are in three folds:

- We construct a novel knowledge-guided RL framework for open attribute value extraction task.
- We provide a benchmark data set for open attribute value extraction task.
- Our method achieves a significantly better performance than the state-of-the-art methods.

2 Overview

Problem Definition

We denote the entity-attribute-value triplet as $\langle e, r, v \rangle$. The goal is to find the attribute value in an incomplete triplet $\langle e, r, ? \rangle$. To achieve this purpose, we pose a question generated with a pre-defined template to search engine to obtain relevant articles. For example, to fill the incomplete triplet $\langle GTX1080, Core\ code, ? \rangle$, we retrieve articles with the query “What is the core code of GTX1080?”.

An information extraction system, such as QANet (Yu et al., 2018), is used to extract a candidate answer with a certain confidence score from

an article. However, due to the inconsistent qualities of the online articles and the inevitable errors caused by the information extraction system, the results of extracted from only one online article is not satisfactory in many cases. Another source of information that can be leveraged help to fulfill such a task is a KG. While it is hard to find out the attribute values for an emerging entity given the existing ones, the KG can serve as the background knowledge about the attributes. We approach the problem using a reinforcement learning framework that is illustrated in the next section.

System Overview

Our procedure is summarized in Figure 1. We use $\langle GTX1180, core\ code, ? \rangle$ as an example for illustration. The query “What is the core code of GTX1180?” is posed to the search engine to obtain a collection of relevant articles by downloading the top M headlines and bodies in the searching page. M is a pre-determined parameter that controls the maximum capacity of the retrieved articles. For each of the retrieved articles, we use an information extraction system to extract a candidate answer. In our example, $RTX2080$ is extracted with a confidence score of 0.30 from the first queried article and $GV104$ is extracted with a confidence score of 0.25 from the second article. Given the first two candidate answers, the RL decides on which answer to pick and whether more articles need to be retrieved.

To make such decisions, in addition to the confidence evidence from the information extraction system, the relevant facts in the KG will be fed into the RL agent to serve as the background knowledge about the attribute. For a triplet $\langle e, r, ? \rangle$, we consider v^r as a reference value with respect to the attribute r if there is a triplet $\langle e', r, v^r \rangle$ and e, e' belong to the same category¹ in the KG. In our example, since $GTX1180$ belongs to the category $NVIDIA\ GPU$, and so does $GTX1080$ and $GTX980$, the reference values are retrieved from the fact that the *core code* of $GTX1080$ is $GP104$ and the *core code* of $GTX980$ is $GM204$. Guided by the KG, the RL agent makes successive evaluations and finally outputs the predicted candidate attribute value via a policy network such as DQN (Mnih et al., 2015).

¹The category information is obtained from concept of CN-DBpedia. The knowledge base contains multi-level hierarchy of categories. We use the lowest-level (most specific) category in the hierarchy to derive the reference values.

3 Reinforcement Learning for Open Attribute Value Extraction

The attribute value extraction task is modeled as a *Markov decision process* (MDP), where the RL agent is actively engaged in the decision making process to maximize the reward, which measures the correctness of the extracted attribute values.

The MDP is modeled as a tuple (S, A, T, R) , where $S = \{s\}$ is the space of all possible real-valued vector states; $A = \{a\}$ is the set of actions; $T(s'|s, a)$ refers to a transition function that maps the domain of state and action to a probability distribution of states; $R(s, a)$ is a reward function that maps the domain of state and action to a real number, which is encoded such that the higher value the better. We describe our RL methodology by illustrating these components as follows.

Action and transition At each decision stage, the agent will observe two candidate answers from two articles and make decisions to answer the two questions: (i) which answer is better out of the two? (ii) should the agent stop at *the current best answer* or continue querying more articles? At the initial decision point, two candidate answers are obtained from two articles simultaneously queried from the web, where we arbitrarily assign one of them to be the current best answer and the other the new candidate answer.

We define the following three actions in A :

1. *Retain*: (i) retain the current best answer and discard the new answer; (ii) query next.
2. *Replace*: (i) replace the current best answer with the new answer; (ii) query next.
3. *Stop*: (i) select the current best answer as the final answer; (ii) stop the query.

At all subsequent decision points, we will retain or replace the *current best answer* and continue comparing with the *new candidate answers* queried from the web until the action is “*Stop*”.

State At each decision point, the state is constructed by concatenating the following three components, where different sources of information are combined.

(1) State variables associated with the confidence scores. The first component is the confidence scores associated with the two candidate answers, which are defined by the information extraction system. We consider this part as the signal

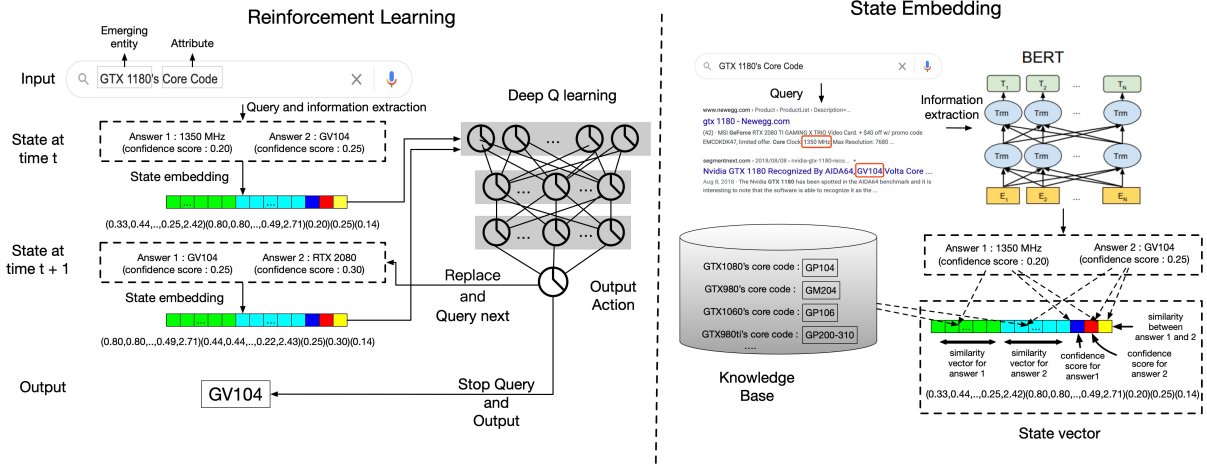


Figure 2: An example of reinforcement learning framework and state embedding. On the left panel, the inputs are the emerging entities and attributes. *The current best answer* and *the next candidate answer* are extracted from the retrieved articles with an information extraction system. The extracted answers are embedded into a state vector from the state embedding process. The state vector is fed into a policy network (DQN). The policy network selects the optimal actions and outputs the current best result. On the right panel, it presents the state embedding process. The extracted candidate answers are embedded into a state vector via similarity metrics and confidence scores.

of the goodness of the extracted answers related to the articles.

(2) State variables informed by the KG. The second component leverages knowledge from the reference values. For a given attribute, we expect the attribute values to be similar to each other in lexical sense. In order to capture such information, we first construct 7 features based on 2 string lexical similarity metrics. For each of the 7 features, we take the average and maximum of the features for each of the two candidate answers as state variables.

String similarity metrics The two string lexical similarity metrics as follows:

$$L_Sim(s_1, s_2) = 1 - \frac{L(s_1, s_2)}{\max(|s_1|, |s_2|)},$$

$$LCS_Sim(s_1, s_2) = \frac{|LCS(s_1, s_2)|}{\max(|s_1|, |s_2|)},$$

where $L(s_1, s_2)$ refers to Levenshtein distance (Levenshtein, 1966). It measures how different two strings are by counting the number of deletions, insertions or substitutions required to transform s_1 into s_2 . $L_Sim(s_1, s_2)$ is known as the Levenshtein similarity and $LCS(s_1, s_2)$ stands for the longest common sub-string of s_1 and s_2 (Gusfield, 1997).

Features based on similarity We define the following 7 features to capture the similarity between two strings from different aspects:

- f_1 : L_Sim between s_1 and s_2 ;
- f_2 : LCS_Sim between s_1 and s_2 ;
- f_3 : L_Sim between s_1 and s_2 with numbers removed from s_1, s_2 ;
- f_4 : LCS_Sim between s_1 and s_2 with numbers removed from s_1, s_2 ;
- f_5 : L_Sim between s_1 and s_2 with s_1, s_2 wildcard masked²;
- f_6 : LCS_Sim between s_1 and s_2 with s_1, s_2 wildcard masked;
- f_7 : The difference in the length of s_1 and s_2 in characters.

Construction of state variables by the KG For each of the 7 features, given the reference values in V^r and the two candidate answers, $answer_1$ and $answer_2$, we form the 28 state variables in this part by taking averages and maximums, which is specified as follows:

$$\left\{ \begin{array}{l} \frac{1}{|V^r|} \sum_{\ell=1}^{|V^r|} f_i(answer_1, v_\ell^r), \\ \max_{\ell=1}^{|V^r|} f_i(answer_1, v_\ell^r), \\ \frac{1}{|V^r|} \sum_{\ell=1}^{|V^r|} f_i(answer_2, v_\ell^r), \\ \max_{\ell=1}^{|V^r|} f_i(answer_2, v_\ell^r), \end{array} \right.$$

for $i = 1, \dots, 7$. This is the part of the state where knowledge from KG is used to inform the decision of the RL agent.

²Wildcard mask means masking the numbers in the string. For example, a string “750 by 1334 Pixels” will be masked to “xxx by xxxx Pixels”.

(3) State variables based on the candidate answers The third component contains the Levenshtein similarity between the two candidate answers. Intuitively, when the confidence scores of both candidate answers are high and they are very similar to each other, then it shows some positive signal for stopping.

The components (1) - (3) are concatenated together to construct the 31-dimensional state vector to carry information from different perspectives.

Reward The reward is set to 0 when the query process is ongoing; only at the final stage when the query is terminated, a nonzero reward is received, which measures the similarity between the final answer and the correct answer.

$$R(s, a) = \begin{cases} L_Sim(\hat{v}, v) & a \text{ is } Stop, \\ 0 & \text{otherwise,} \end{cases}$$

where \hat{v} is the selected best attribute value and v is the true attribute value.

Method Since the state defined in our framework is from a continuous space, we adopt a deep Q-network (DQN) to approximate $Q(s, a)$ with a deep neural network denoted by $Q(s, a; \theta)$. Specifically, we parameterize an approximate value function $Q(s, a; \theta)$ using a three-layer deep neural network. The network takes the continuous 31-dimensional state vector s as input and predict $Q(s, a)$. We use the rectified linear unit (ReLU) activation functions in the hidden layers. The architecture is illustrated in Figure 2.

Algorithm 1 provides complete details of our MDP framework for the DQN training phase.

4 Experiments

In this section, we compare our proposed RL framework to the state-of-the-art extraction-based baselines, demonstrating its robustness and ability to obtain accurate answers for missing attribute values. Our codes are publicly available online.³

4.1 Data

The dataset is generated from existing triplets using the largest public Chinese knowledge base, CN-DBpedia (Xu et al., 2017), with a corresponding taxonomy CN-Probase.⁴ Specially, the number of

³<https://github.com/yeliu0930/Knowledge-guided-Open-Attribute-Value-Extraction-with-Reinforcement-Learning>

⁴As far as we know, there is no public benchmark dataset suitable for our open attribute value extraction task with labeled values.

Algorithm 1 The full details of our training Phase for the DQN agent with ϵ -greedy exploration.

```

1: Initialize a set of training triplets
    $x_i = \langle e_i, r_i, v_i \rangle \in X$ 
2: Initialize parameters  $\theta$  randomly
3: Initialize replay memory  $\mathcal{D}$ 
4: for  $x_i \in X$  do
5:   Download M articles by searching with
     query “[ $e_i$ ] $'s[r_i]$ ”
6:   Queue the downloaded articles in  $C_i$ 
7:   Identify reference values from the KG and
     save them in  $V_i^r$ 
8:   for epoch = 1, ..., E do
9:     for  $i = 1, \dots, |X|$  do
10:      Pop the first two articles in  $C_i$  and obtain
         $answer_1$  with  $confidence_1$  and  $answer_2$ 
        with  $confidence_2$ 
11:      Form the state  $s_1$  given  $answer_1$ ,  $confidence_1$ ,
         $answer_2$ ,  $confidence_2$ ,  $V_i^r$ 
12:      for  $t = 1, \dots, M - 1$  do
13:        With probability  $1 - \epsilon$  select  $a_t =$ 
         $\text{argmax}_a Q(s_t, a; \theta)$  otherwise select
         $a_t$  randomly
14:        if  $a_t$  is not “Stop” then
15:           $r_t \leftarrow 0$ 
16:          Pop next article from  $C_i$  and obtain
             $answer_{new}$  with  $confidence_{new}$ 
17:          if  $a_t$  is “Retain” then
18:             $answer_2 \leftarrow answer_{new}$ 
19:             $confidence_2 \leftarrow confidence_{new}$ 
20:          if  $a_t$  is “Replace” then
21:             $answer_1 \leftarrow answer_2$ 
22:             $confidence_1 \leftarrow confidence_2$ 
23:             $answer_2 \leftarrow answer_{new}$ 
24:             $confidence_2 \leftarrow confidence_{new}$ 
25:          Form a new state  $s_{t+1}$ 
26:        else
27:           $\hat{v}_i \leftarrow answer_1$ 
28:           $r_t = L\_Sim(\hat{v}_i, v_i)$ 
29:           $s_{t+1} \leftarrow NULL$ 
30:        Store transition  $(s_t, a_t, r_t, s_{t+1})$  in  $\mathcal{D}$ 
31:        Sample random mini batch of transi-
        tions  $(s_t, a_t, r_t, s_{t+1})$  from  $\mathcal{D}$ 
32:         $y_t = r_t$  if  $a_t$  is “Stop”
        else  $r_t + \gamma \max_{a'} Q(s_{t+1}, a'; \theta)$ 
33:        Update parameter  $\theta$  on the loss  $\mathcal{L}(\theta) =$ 
         $(y_t - Q(s_t, a_t; \theta))^2$ 
34:        if  $a_t$  is “Stop” then break

```

training triplets is 1022. The selected entities in the experiment are from four different fields, including GPU, game, movie and phone. The testing data contains 75 triplets for each field, hence the total number of triplets in the testing is 300. For each triplet in the training and testing data, we download articles from top $M = 10$ links obtained from the Baidu search engine. The CN-DBpedia is used as our external KG with the triplets in training and testing masked and to provide reference values.

4.2 Reinforcement Learning Implementation

In the RL setting, we use DQN to train the policy. Specifically, the DQN contains three layers of multilayer perceptron (MLP). The dimensions of hidden layers in MLP are chosen as 10 and 5 respectively. The dimension for the output is 3 which represents the three actions. In our experiments, the DQN model is trained 100 epochs where each epoch contains 1,000 transitions. We use a decreasing learning rate with epochs during the training process⁵. The ϵ in ϵ -greedy exploration is annealed from 1 to 0.02 over 10,000 transitions. The replay memory \mathcal{D} is of size 10,000. We deploy our RL model in RLlib (Liang et al., 2017) for efficiently distributed computation.

4.3 Information Extraction System

Different information extraction methods are implemented during the experiment. Sequence labeling methods including Bi-LSTM labeling, Bi-LSTM-CRF labeling (Huang et al., 2015), CNN labeling (Collobert et al., 2011), CNN-att-CRF labeling (Tan et al., 2018), and OpenTag labeling (Zheng et al., 2018) are used. We also consider three *Machine Comprehension* (MC) models, including BiDAF (Seo et al., 2016), QANet and a Bert-based model (Devlin et al., 2018). The SQuAD-like Chinese open-domain MC datasets including WebQA (Li et al., 2016b) and CIPS-SOGO factoid question-answering subtask dataset⁶ are used to train the extraction systems with distant supervision.

4.4 Competitors

We experimented with three traditional aggregation methods and four variants of RL agents as the competitors during the experiment.

Traditional Aggregation Methods

⁵The learning rate schedule is set to $[[0, 0.05], [20, 0.01], [30, 0.005], [50, 0.001]]$.

⁶http://task.www.sogou.com/cips-sogou_qa/

1. Random choice (*Random*): We randomly select an article out of M articles and extract an answer from it with the information extraction system as the final answer.

2. First article (*First*): We use the answer extracted from the article that ranked first in the search engines.

3. Majority aggregation (*Majority*): We use a majority vote strategy over all the extracted answers.

4. Confidence aggregation (*Confidence*): The answer with the highest confidence score out is chosen as the final answer. This aggregation method is only feasible when each candidate answer is associated with a confidence score.

Variations of the RL framework

1. *RL-NK*: (No KG included) The RL agent do not leverage the information from KG. The KG-dependent part (i.e. the component (2) in state construction) is omitted from the state.

2. *RL-NR*: (No retain or replace actions) The only action in the RL framework is *Stop*. The final answer is the one with the highest confidence score among candidate answers seen before stop.

3. *RL-NS*: (No stop action) The RL agent do not make *Stop* decisions. All of the M extracted candidate answers are compared.

4. *RL-KG*: Our proposed RL framework.

Since the sequence labeling methods cannot provide valid confidence scores associated with the candidate answers, the answers extracted with these methods are aggregated using *Random* and *Majority* strategies. For the MC models, we implemented all the aggregation strategies including our RL-based methods.

4.5 Results

Our evaluating metric is the Levenshtein similarity between the final answer and the ground truth, which ranges from 0 to 1 and higher score represents better performance. The results are summarized in Table 1 when different information extraction systems are combined with different aggregation strategies. The results are evaluated separately under each field and the combined results are also reported in the tables. All results reported are averaged over 3 independent runs. The oracle performances are provided to differentiate the error incurred by imperfect decisions and the inherent errors caused by the information extraction system.

	Evaluating Dataset				
	GPU	Games	Movie	Phone	All
Baselines: Sequence labeling methods with traditional aggregations					
<i>Random</i> (OpenTag)	0.222	0.220	0.238	0.244	0.231
<i>Majority</i> (OpenTag)	0.282	0.334	0.321	0.320	0.314
<i>Random</i> (Bi-LSTM)	0.291	0.322	0.184	0.418	0.304
<i>Majority</i> (Bi-LSTM)	0.307	0.334	0.194	0.462	0.324
<i>Random</i> (Bi-LSTM-CRF)	0.349	0.273	0.287	0.336	0.311
<i>Majority</i> (Bi-LSTM-CRF)	0.517	0.388	0.360	0.494	0.440
<i>Random</i> (CNN)	0.409	0.333	0.356	0.381	0.370
<i>Majority</i> (CNN)	0.534	0.419	0.430	0.508	0.473
<i>Random</i> (CNN-att-CRF)	0.399	0.272	0.304	0.310	0.321
<i>Majority</i> (CNN-att-CRF)	0.626	0.413	0.481	0.506	0.507
BiDAF with traditional aggregations, RL methods and Oracle strategy.					
<i>Random</i> (BiDAF)	0.259	0.155	0.267	0.223	0.226
<i>First</i> (BiDAF)	0.451	0.498	0.533	0.632	0.528
<i>Majority</i> (BiDAF)	0.488	0.321	0.539	0.415	0.441
<i>Confidence</i> (BiDAF)	0.799	0.488	0.645	0.560	0.623
<i>RL-NK</i> (BiDAF)	0.679	0.602	0.609	0.658	0.637
<i>RL-NR</i> (BiDAF)	0.751	0.655	0.622	0.644	0.668
<i>RL-NS</i> (BiDAF)	0.759	0.732	0.673	0.680	0.711
<i>RL-KG</i> (BiDAF)	0.786	0.692	0.686	0.739	0.726
<i>Oracle</i> (BiDAF)	0.902	0.793	0.846	0.812	0.838
QANet with traditional aggregations, RL methods and Oracle strategy					
<i>Random</i> (QANet)	0.261	0.167	0.259	0.236	0.230
<i>First</i> (QANet)	0.507	0.533	0.531	0.675	0.561
<i>Majority</i> (QANet)	0.484	0.325	0.500	0.469	0.444
<i>Confidence</i> (QANet)	0.691	0.493	0.689	0.546	0.605
<i>RL-NK</i> (QANet)	0.640	0.592	0.596	0.645	0.618
<i>RL-NR</i> (QANet)	0.687	0.717	0.549	0.631	0.646
<i>RL-NS</i> (QANet)	0.801	0.695	0.686	0.771	0.738
<i>RL-KG</i> (QANet)	0.786	0.687	0.731	0.790	0.749
<i>Oracle</i> (QANet)	0.932	0.840	0.878	0.868	0.880
BERT with traditional aggregations, RL methods and Oracle strategy					
<i>Random</i> (BERT)	0.374	0.234	0.361	0.287	0.314
<i>First</i> (BERT)	0.507	0.533	0.531	0.675	0.561
<i>Majority</i> (BERT)	0.620	0.438	0.626	0.530	0.553
<i>Confidence</i> (BERT)	0.727	0.600	0.552	0.540	0.605
<i>RL-NK</i> (BERT)	0.716	0.682	0.565	0.687	0.662
<i>RL-NR</i> (BERT)	0.775	0.652	0.707	0.723	0.714
<i>RL-NS</i> (BERT)	0.773	0.673	0.769	0.831	0.762
<i>RL-KG</i> (BERT)	0.817	0.637	0.777	0.837	0.767
<i>Oracle</i> (BERT)	0.925	0.857	0.887	0.909	0.895

Table 1: Accuracy of the baseline methods and our proposed methods. Bold indicates best baseline performances with a sequence labeling methods and best results achieved with BiDAF/QANet/BERT. The Oracle performance shows the best possible performance when perfect decisions are made. **Our proposed RL-KG improves the extraction performances substantially.**

From Table 1, we have the following observations. First, the RL based methods outperform all the competing baseline methods. By adopting the RL framework instead of traditional aggregation methods, the accuracies are boosted substantially. It demonstrate the effectiveness of the RL framework. Second, compared to the RL framework without the guide of KG (*RL-NK*), our proposed *RL-KG* framework achieves significantly better results. This suggests that the KG does provide valuable information in the task of the attribute value extraction. Third, the *RL-KG* framework outperforms all the other variants of the RL framework. It

shows that considering answer selection and stopping decisions at the same time achieves the best performances.

We also conduct an experiment to see how our method performs when the KG is not able to provide information for some triplets, which is a common situation in reality. During the experiment, we randomly set the reference values for 0 – 100 (incremented by 10) percent of triplets in the training and testing as empty. For those triplets that do not have reference values, the state variables associated with the KG (i.e. the component (2) in our state construction) are set to 0. Figure 3 displays how the performances change when information from KG is leveraged with different levels of frequency. It can be seen that for all the three information extraction models, the performances are getting better as the KG is used at higher frequencies.




Figure 3: Extraction accuracy when KG is used with different levels of frequency. **Better extraction results are achieved when KG is used with higher frequencies.**

4.6 Case Study

By incorporating information from KG, RL agent is able to rule out some unreasonable answers and boost the extraction accuracy.

In Table 2, we present some cases where the trained RL agent helps to correct the information extraction errors. More details for the first two examples are included in Table 3. For the first example in Table 3, the emerging entity is *Founder r680-470* and the attribute of interest is *operating system*. The reference values retrieved from the KG include values like *IOS*, *DOS*, *Andriod*, *EMUI*. The trained RL agent stops after querying the six raw corpus and selects the answer *DOS* over the previous candidate values, which is exactly one of the reference values. In the second example in Table 3, we are interested in the *release date* for the emerging entity *super puzzles game*. The reference val-

Entity	Attribute	Truth	<i>RL-KG(BERT)</i>	<i>Confidence(BERT)</i>
Founder r680-470	Operating system	DOS	DOS	T2410
Super puzzles game	Release date	12/25/2007	12/25/2007	43.75
Madea’s Witness Protection	Rating system	USA:PG-13	pg-13	6.4
Tequila Works, S.L	Number of endings	2	2	ConsoleGame

Table 2: Case studies (translated)

Example 1	
< <i>e, r, ?</i> >	<Founder r680-470, Operating system, ? >
...	... ✗
Raw corpus 4	Intel Pentium dual-core T2410 is an entry level processor based on the Merom-2M... ✗
Raw corpus 5	...Founder Q680 high shot instrument scanner 5 million pixels A4 on... ✗
Raw corpus 6	...The latest R680 series laptop, Operating system DOS with ... ✓
...	...
Reference values	<..IOS, DOS, Andriod, EMUI..>
Example 2	
< <i>e, r, ?</i> >	<Super puzzles game, Release date, ? >
Raw corpus 1	...announced release date of super puzzles is 12/25/2007 . It was published by... ✓
Raw corpus 2	...Super Puzzle HD v2.1.2 - Size: 43.75 MB...
...	...
Reference values	<..02/24, 03/27/2018, 05/07/2016..>

Table 3: Detailed examples where the trained RL agent helps to select the candidate answers from raw corpus. Blue words represent the extracted answers from BERT. The check-marks denote the selected answers.

ues are dates like *02/24, 03/27/2018, 05/07/2016*, etc. The trained RL agent stops after one step, and outputs the candidate answer *12/25/2007*. By leveraging the reference values, our proposed method demonstrated its advantages in answer selection.

5 Related Work

5.1 Machine reading comprehension

Machine reading comprehension (MRC) and automated question (QA) answering are important and longstanding topic in NLP research due to its huge potentials in wide variety of applications. An end-to-end MRC QA models are expected to have the ability to read a piece of text and then answer questions about it. Significant progress has been made with the machine reading and QA task in recent years. Some notable works include BiDAF (Seo et al., 2016), SAN (Liu et al., 2017), QANet, ALBERT (Lan et al., 2019).

Our proposed framework can also be regarded an end-to-end MRC QA model that is built on top of an existing MRC QA model, which is used as the information extraction system in our extraction process. Different from most of the previous works, our focus is to enhance the performance of an ex-

isting model by utilizing external information from KG and by acquiring more articles when the agent does not feel confident about the extracted answer.

5.2 Open-world knowledge graph completion

Attribute value extraction under the open world assumption has received many attentions in NLP community recently. There has been quite a few works on open attribute value extraction. Open-Tag (Zheng et al., 2018) formalized the extraction problem as a sequence tagging task and proposed an end-to-end framework for open attribute value extraction. The open-world KGC (Shi and Weninger, 2018) used a complex relationship dependent content masking architecture to mitigate the presence of noisy text descriptions and extract the attribute value from the denoised text. TXtract (Karamanolakis et al., 2020) incorporated the categorical structure into the value tagging system. However these methods suffer from irrelevant articles and is not able to filter out noisy answers.

5.3 NLP with reinforcement learning

RL (Sutton et al., 1998) is a framework that enables agents to reason about sequential decision making

as an optimization process. It has been widely applied in NLP tasks, including article summarization (Paulus et al., 2017; Li et al., 2018; Celikyilmaz et al., 2018), dialogue generation (Li et al., 2016a; Serban et al., 2017; Li et al., 2019), and question answering (Xiong et al., 2017; Wang et al., 2018; Das et al., 2019) and so on. To the best of our knowledge, we are the first to integrate information from KG into a RL framework to fulfill the attribute extraction task.

6 Conclusion and discussion

This paper presents a novel RL framework to perform open attribute value extraction. Through a set of experiments, we observe that the most of the computation cost is incurred by training the information extraction system. The remaining computation cost from RL framework is comparably small during both the training and the prediction process. Specifically, during our experiments, we trained a three-layer deep neural network model, which has much fewer parameters compared to the information extraction system. The proposed RL method demonstrates promising performance, where the KG showed its ability to provide guidance in open attribute extraction task. Our framework also contributes to areas of knowledge graph completion and automatic question-answering for attribute values.

KG has huge potential to provide rich background information in many NLP applications. Our solution for attribute value extraction can be extended to other NLP tasks. A potential attempt might be to use KG to design the reward in the RL framework to provide weak supervision. We leave this as our future work.

References

Asli Celikyilmaz, Antoine Bosselut, Xiaodong He, and Yejin Choi. 2018. Deep communicating agents for abstractive summarization. *arXiv preprint arXiv:1803.10357*.

Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of machine learning research*, 12(Aug):2493–2537.

Rajarshi Das, Shehzaad Dhuliawala, Manzil Zaheer, and Andrew McCallum. 2019. Multi-step retriever-reader interaction for scalable open-domain question answering. *arXiv preprint arXiv:1905.05733*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Dan Gusfield. 1997. Algorithms on stings, trees, and sequences: Computer science and computational biology. *Acm Sigact News*, 28(4):41–60.

Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional lstm-crf models for sequence tagging. *arXiv preprint arXiv:1508.01991*.

Giannis Karamanolakis, Jun Ma, and Xin Luna Dong. 2020. Textract: Taxonomy-aware knowledge extraction for thousands of product categories. *arXiv preprint arXiv:2004.13852*.

Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.

Vladimir I Levenshtein. 1966. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*, volume 10, pages 707–710.

Jiwei Li, Will Monroe, Alan Ritter, Michel Galley, Jianfeng Gao, and Dan Jurafsky. 2016a. Deep reinforcement learning for dialogue generation. *arXiv preprint arXiv:1606.01541*.

Peng Li, Wei Li, Zhengyan He, Xuguang Wang, Ying Cao, Jie Zhou, and Wei Xu. 2016b. Dataset and neural recurrent sequence labeling model for open-domain factoid question answering. *arXiv preprint arXiv:1607.06275*.

Piji Li, Lidong Bing, and Wai Lam. 2018. Actor-critic based training framework for abstractive summarization. *arXiv preprint arXiv:1803.11070*.

Ziming Li, Julia Kiseleva, and Maarten de Rijke. 2019. Dialogue generation: From imitation learning to inverse reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6722–6729.

Eric Liang, Richard Liaw, Philipp Moritz, Robert Nishihara, Roy Fox, Ken Goldberg, Joseph E Gonzalez, Michael I Jordan, and Ion Stoica. 2017. Rllib: Abstractions for distributed reinforcement learning. *arXiv preprint arXiv:1712.09381*.

Xiaodong Liu, Yelong Shen, Kevin Duh, and Jianfeng Gao. 2017. Stochastic answer networks for machine reading comprehension. *arXiv preprint arXiv:1712.03556*.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533.

- Karthik Narasimhan, Adam Yala, and Regina Barzilay. 2016. Improving information extraction by acquiring external evidence with reinforcement learning. *arXiv preprint arXiv:1603.07954*.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2017. A deep reinforced model for abstractive summarization. *arXiv preprint arXiv:1705.04304*.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2016. Bidirectional attention flow for machine comprehension. *arXiv preprint arXiv:1611.01603*.
- Iulian V Serban, Chinnadhurai Sankar, Mathieu Germain, Saizheng Zhang, Zhouhan Lin, Sandeep Subramanian, Taesup Kim, Michael Pieper, Sarath Chandar, Nan Rosemary Ke, et al. 2017. A deep reinforcement learning chatbot. *arXiv preprint arXiv:1709.02349*.
- Baoxu Shi and Tim Wenginger. 2018. Open-world knowledge graph completion. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Richard S Sutton, Andrew G Barto, et al. 1998. *Introduction to reinforcement learning*, volume 135. MIT press Cambridge.
- Zhixing Tan, Mingxuan Wang, Jun Xie, Yidong Chen, and Xiaodong Shi. 2018. Deep semantic role labeling with self-attention. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Shuohang Wang, Mo Yu, Xiaoxiao Guo, Zhiguo Wang, Tim Klinger, Wei Zhang, Shiyu Chang, Gerry Tesauro, Bowen Zhou, and Jing Jiang. 2018. R 3: Reinforced ranker-reader for open-domain question answering. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Caiming Xiong, Victor Zhong, and Richard Socher. 2017. Dcn+: Mixed objective and deep residual coattention for question answering. *arXiv preprint arXiv:1711.00106*.
- Bo Xu, Yong Xu, Jiaqing Liang, Chenhao Xie, Bin Liang, Wanyun Cui, and Yanghua Xiao. 2017. Cndbpedia: A never-ending chinese knowledge extraction system. In *International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems*, pages 428–438. Springer.
- Adams Wei Yu, David Dohan, Minh-Thang Luong, Rui Zhao, Kai Chen, Mohammad Norouzi, and Quoc V Le. 2018. Qanet: Combining local convolution with global self-attention for reading comprehension. *arXiv preprint arXiv:1804.09541*.
- Guineng Zheng, Subhabrata Mukherjee, Xin Luna Dong, and Feifei Li. 2018. Opentag: Open attribute value extraction from product profiles. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1049–1058.