

RNNLOGIC: LEARNING LOGIC RULES FOR REASONING ON KNOWLEDGE GRAPHS

Meng Qu^{*1,2}, Junkun Chen^{*3}, Louis-Pascal Xhonneux^{1,2}, Yoshua Bengio^{1,2,5}, Jian Tang^{1,4,5}

¹Mila - Québec AI Institute

²Université de Montréal

³Tsinghua University

⁴HEC Montréal

⁵Canadian Institute for Advanced Research (CIFAR)

ABSTRACT

This paper studies learning logic rules for reasoning on knowledge graphs. Logic rules provide interpretable explanations when used for prediction as well as being able to generalize to other tasks, and hence are critical to learn. Existing methods either suffer from the problem of searching in a large search space (e.g., neural logic programming) or ineffective optimization due to sparse rewards (e.g., techniques based on reinforcement learning). To address these limitations, this paper proposes a probabilistic model called RNNLogic. RNNLogic treats logic rules as a latent variable, and simultaneously trains a rule generator as well as a reasoning predictor with logic rules. We develop an EM-based algorithm for optimization. In each iteration, the reasoning predictor is first updated to explore some generated logic rules for reasoning. Then in the E-step, we select a set of high-quality rules from all generated rules with both the rule generator and reasoning predictor via posterior inference; and in the M-step, the rule generator is updated with the rules selected in the E-step. Experiments on four datasets prove the effectiveness of RNNLogic.

1 INTRODUCTION

Knowledge graphs are collections of real-world facts, which are useful in various applications. Each fact is typically specified as a triplet (h, r, t) or equivalently $r(h, t)$, meaning entity h has relation r with entity t . For example, *Bill Gates* is the *Co-founder* of *Microsoft*. As it is impossible to collect all facts, knowledge graphs are incomplete. Therefore, a fundamental problem on knowledge graphs is to predict missing facts by reasoning with existing ones, a.k.a. knowledge graph reasoning.

This paper studies learning logic rules for reasoning on knowledge graphs. For example, one may extract a rule $\forall X, Y, Z \text{ hobby}(X, Y) \leftarrow \text{friend}(X, Z) \wedge \text{hobby}(Z, Y)$, meaning that if Z is a friend of X and Z has hobby Y , then Y is also likely the hobby of X . Then the rule can be applied to infer new hobbies of people. Such logic rules are able to improve interpretability and precision of reasoning (Qu & Tang, 2019; Zhang et al., 2020). Moreover, logic rules can also be reused and generalized to other domains and data (Teru & Hamilton, 2020). However, due to the large search space, inferring high-quality logic rules for reasoning on knowledge graphs is a challenging task.

Indeed, a variety of methods have been proposed for learning logic rules from knowledge graphs. Most traditional methods such as path ranking (Lao & Cohen, 2010) and Markov logic networks (Richardson & Domingos, 2006) enumerate relational paths on graphs as candidate logic rules, and then learn a weight for each rule as an assessment of rule qualities. There are also some recent methods based on neural logic programming (Yang et al., 2017) and neural theorem provers (Rocktäschel & Riedel, 2017), which are able to learn logic rules and their weights simultaneously in a differentiable way. Though empirically effective for prediction, the search space of these methods is exponentially large, making it hard to identify high-quality logic rules. Besides, some recent efforts (Xiong et al., 2017) formulate the problem as a sequential decision making process, and use reinforcement learning to search for logic rules, which significantly reduces search complexity. However, due to the large action space and sparse reward in training, the performance of these methods is not yet satisfying.

^{*}Equal contribution.

In this paper, we propose a principled probabilistic approach called RNNLogic which overcomes the above limitations. Our approach consists of a rule generator as well as a reasoning predictor with logic rules, which are simultaneously trained to enhance each other. The rule generator provides logic rules which are used by the reasoning predictor for reasoning, while the reasoning predictor provides effective reward to train the rule generator, which helps significantly reduce the search space. Specifically, for each query-answer pair, e.g., $\mathbf{q} = (h, r, ?)$ and $\mathbf{a} = t$, we model the probability of the answer conditioned on the query and existing knowledge graph \mathcal{G} , i.e., $p(\mathbf{a}|\mathcal{G}, \mathbf{q})$, where a set of logic rules \mathbf{z} ¹ is treated as a latent variable. The rule generator defines a prior distribution over logic rules for each query, i.e., $p(\mathbf{z}|\mathbf{q})$, which is parameterized by a recurrent neural network. The reasoning predictor computes the likelihood of the answer conditioned on the logic rules and the existing knowledge graph \mathcal{G} , i.e., $p(\mathbf{a}|\mathcal{G}, \mathbf{q}, \mathbf{z})$. At each training iteration, we first sample a few logic rules from the rule generator, and further update the reasoning predictor to try out these rules for prediction. Then an EM algorithm (Neal & Hinton, 1998) is used to optimize the rule generator. In the E-step, a set of high-quality logic rules are selected from all the generated rules according to their posterior probabilities. In the M-step, the rule generator is updated to imitate the high-quality rules selected in the E-step. Extensive experimental results show that RNNLogic outperforms state-of-the-art methods for knowledge graph reasoning². Besides, RNNLogic is able to generate high-quality logic rules.

2 RELATED WORK

Our work is related to existing efforts on learning logic rules for knowledge graph reasoning. Most traditional methods enumerate relational paths between query entities and answer entities as candidate logic rules, and further learn a scalar weight for each rule to assess the quality. Representative methods include Markov logic networks (Kok & Domingos, 2005; Richardson & Domingos, 2006; Khot et al., 2011), relational dependency networks (Neville & Jensen, 2007; Natarajan et al., 2010), rule mining algorithms (Galárraga et al., 2013; Meilicke et al., 2019), path ranking (Lao & Cohen, 2010; Lao et al., 2011) and probabilistic personalized page rank (ProPPR) algorithms (Wang et al., 2013; 2014a;b). Some recent methods extend the idea by simultaneously learning logic rules and the weights in a differentiable way, and most of them are based on neural logic programming (Rocktäschel & Riedel, 2017; Yang et al., 2017; Cohen et al., 2018; Sadeghian et al., 2019; Yang & Song, 2020) or neural theorem provers (Rocktäschel & Riedel, 2017; Minervini et al., 2020). These methods and our approach are similar in spirit, as they are all able to learn the weights of logic rules efficiently. However, these existing methods try to simultaneously learn logic rules and their weights, which is nontrivial in terms of optimization. The main innovation of our approach is to separate rule generation and rule weight learning by introducing a rule generator and a reasoning predictor respectively, which can mutually enhance each other. The rule generator generates a few high-quality logic rules, and the reasoning predictor only focuses on learning the weights of such high-quality rules, which significantly reduces the search space and leads to better reasoning results. Meanwhile, the reasoning predictor can in turn help identify some useful logic rules to improve the rule generator.

The other kind of rule learning method is based on reinforcement learning. The general idea is to train pathfinding agents, which search for reasoning paths in knowledge graphs to answer questions, and then extract logic rules from reasoning paths (Xiong et al., 2017; Chen et al., 2018; Das et al., 2018; Lin et al., 2018; Shen et al., 2018). However, training effective pathfinding agents is highly challenging, as the reward signal (i.e., whether a path ends at the correct answer) can be extremely sparse. Although some studies (Lin et al., 2018) try to get better reward by using embedding-based methods for reward shaping, the performance is still worse than most embedding-based methods. In our approach, the rule generator has a similar role to those pathfinding agents. The major difference is that we simultaneously train the rule generator and a reasoning predictor with logic rules, which mutually enhance each other. The reasoning predictor provides effective reward for training the rule generator, and the rule generator offers high-quality rules to improve the reasoning predictor.

Our work is also related to knowledge graph embedding, which solves knowledge graph reasoning by learning entity and relation embeddings in latent spaces (Bordes et al., 2013; Wang et al., 2014c; Yang et al., 2015; Nickel et al., 2016; Trouillon et al., 2016; Cai & Wang, 2018; Dettmers et al., 2018; Balazevic et al., 2019; Sun et al., 2019). With proper architectures, these methods are able to learn

¹More precisely, \mathbf{z} is a multiset. In this paper, we use “set” to refer to “multiset” for conciseness.

²The codes of RNNLogic are available: <https://github.com/DeepGraphLearning/RNNLogic>

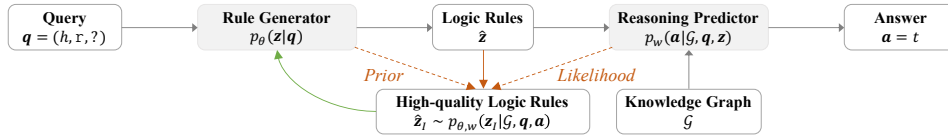


Figure 1: RNNLogic consists of a rule generator p_θ and a reasoning predictor p_w . Given a query, the rule generator generates logic rules for the reasoning predictor. The reasoning predictor takes the generated rules as input, and reasons on a knowledge graph to predict the answer. RNNLogic is optimized with an EM-based algorithm. In each iteration, the rule generator produces some logic rules, and we update the reasoning predictor to explore these rules for reasoning. Then in the **E-step**, a set of high-quality rules are identified from all generated rules via posterior inference. Finally in the **M-step**, the rule generator is updated to be consistent with the high-quality rules identified in E-step.

some simple logic rules. For example, TransE (Bordes et al., 2013) can learn some composition rules. RotatE (Sun et al., 2019) can mine some composition rules, symmetric rules and inverse rules. However, these methods can only find some simple rules in an implicit way. In contrast, our approach explicitly trains a rule generator, which is able to generate more complicated logic rules.

There are some works studying boosting rule-based models (Goldberg & Eckstein, 2010; Eckstein et al., 2017), where they dynamically add new rules according to the rule weights learned so far. These methods have been proven effective in binary classification and regression. Compared with them, our approach shares similar ideas, as we dynamically update the rule generator with the feedback from the reasoning predictor, but we focus on a different task, i.e., reasoning on knowledge graphs.

3 MODEL

In this section, we introduce the proposed approach RNNLogic which learns logic rules for knowledge graph reasoning. We first formally define knowledge graph reasoning and logic rules.

Knowledge Graph Reasoning. Let $p_{\text{data}}(\mathcal{G}, \mathbf{q}, \mathbf{a})$ be a training data distribution, where \mathcal{G} is a background knowledge graph characterized by a set of (h, r, t) -triplets which we may also write as $\mathcal{r}(h, t)$, $\mathbf{q} = (h, r, ?)$ is a query, and $\mathbf{a} = t$ is the answer. Given \mathcal{G} and the query \mathbf{q} , the goal is to predict the correct answer \mathbf{a} . More formally, we aim to model the probabilistic distribution $p(\mathbf{a}|\mathcal{G}, \mathbf{q})$.

Logic Rule. We perform knowledge graph reasoning by learning logic rules, where logic rules in this paper have the conjunctive form $\forall\{X_i\}_{i=0}^l \mathcal{r}(X_0, X_l) \leftarrow \mathcal{r}_1(X_0, X_1) \wedge \dots \wedge \mathcal{r}_l(X_{l-1}, X_l)$ with l being the rule length. This syntactic structure naturally captures composition, and can easily express other common logic rules such as symmetric or inverse rules. For example, let \mathcal{r}^{-1} denote the inverse relation of relation \mathcal{r} , then each symmetric rule can be expressed as $\forall\{X, Y\} \mathcal{r}(X, Y) \leftarrow \mathcal{r}^{-1}(X, Y)$.

In RNNLogic, we treat a set of logic rules which could explain a query as a latent variable we have to infer. To do this, we introduce a rule generator and a reasoning predictor using logic rules. Given a query, the rule generator employs a recurrent neural network to generate a set of logic rules, which are given to the reasoning predictor for prediction. We optimize RNNLogic with an EM-based algorithm. In each iteration, we start with updating the reasoning predictor to try out some logic rules generated by the rule generator. Then in the E-step, we identify a set of high-quality rules from all generated rules via posterior inference, with the prior from the rule generator and likelihood from the reasoning predictor. Finally in the M-step, the rule generator is updated with the identified high-quality rules.

3.1 PROBABILISTIC FORMALIZATION

We start by formalizing knowledge graph reasoning in a probabilistic way, where a *set of logic rules* \mathbf{z} is treated as a latent variable. The target distribution $p(\mathbf{a}|\mathcal{G}, \mathbf{q})$ is jointly modeled by a rule generator and a reasoning predictor. The rule generator p_θ defines a prior over a set of latent rules \mathbf{z} conditioned on a query \mathbf{q} , while the reasoning predictor p_w gives the likelihood of the answer \mathbf{a} conditioned on latent rules \mathbf{z} , the query \mathbf{q} , and the knowledge graph \mathcal{G} . Thus $p(\mathbf{a}|\mathcal{G}, \mathbf{q})$ can be computed as below:

$$p_{w,\theta}(\mathbf{a}|\mathcal{G}, \mathbf{q}) = \sum_{\mathbf{z}} p_w(\mathbf{a}|\mathcal{G}, \mathbf{q}, \mathbf{z})p_\theta(\mathbf{z}|\mathbf{q}) = \mathbb{E}_{p_\theta(\mathbf{z}|\mathbf{q})}[p_w(\mathbf{a}|\mathcal{G}, \mathbf{q}, \mathbf{z})]. \quad (1)$$

The goal is to jointly train the rule generator and reasoning predictor to maximize the likelihood of training data. Formally, the objective function is presented as below:

$$\max_{\theta, w} \mathcal{O}(\theta, w) = \mathbb{E}_{(\mathcal{G}, \mathbf{q}, \mathbf{a}) \sim p_{\text{data}}} [\log p_{w, \theta}(\mathbf{a} | \mathcal{G}, \mathbf{q})] = \mathbb{E}_{(\mathcal{G}, \mathbf{q}, \mathbf{a}) \sim p_{\text{data}}} [\log \mathbb{E}_{p_{\theta}(\mathbf{z} | \mathbf{q})} [p_w(\mathbf{a} | \mathcal{G}, \mathbf{q}, \mathbf{z})]]. \quad (2)$$

3.2 PARAMETERIZATION

Rule Generator. The rule generator defines the distribution $p_{\theta}(\mathbf{z} | \mathbf{q})$. For a query \mathbf{q} , the rule generator aims at generating a set of latent logic rules \mathbf{z} for answering the query.

Formally, given a query $\mathbf{q} = (h, \mathfrak{r}, ?)$, we generate compositional logic rules by only considering the query relation \mathfrak{r} without the query entity h , which allows the generated rules to generalize across entities. For each compositional rule in the abbreviation form $\mathfrak{r} \leftarrow \mathfrak{r}_1 \wedge \dots \wedge \mathfrak{r}_l$, it can be viewed as a sequence of relations $[\mathfrak{r}, \mathfrak{r}_1, \mathfrak{r}_2 \dots \mathfrak{r}_l, \mathfrak{r}_{\text{END}}]$, where \mathfrak{r} is the query relation or the rule head, $\{\mathfrak{r}_i\}_{i=1}^l$ are the body of the rule, and $\mathfrak{r}_{\text{END}}$ is a special relation indicating the end of the relation sequence.

Such relation sequences can be effectively modeled by recurrent neural networks (Hochreiter & Schmidhuber, 1997), and thus we introduce RNN_{θ} to parameterize the rule generator. Given a query relation \mathfrak{r} , RNN_{θ} sequentially generates each relation in the body of a rule, until it reaches the ending relation $\mathfrak{r}_{\text{END}}$. In this process, the probabilities of generated rules are simultaneously computed. With such rule probabilities, we define the distribution over a set of rules \mathbf{z} as a multinomial distribution:

$$p_{\theta}(\mathbf{z} | \mathbf{q}) = \text{Mu}(\mathbf{z} | N, \text{RNN}_{\theta}(\cdot | \mathfrak{r})), \quad (3)$$

where Mu stands for multinomial distributions, N is a hyperparameter for the size of the set \mathbf{z} , and $\text{RNN}_{\theta}(\cdot | \mathfrak{r})$ defines a distribution over compositional rules with rule head being \mathfrak{r} . The generative process of a rule set \mathbf{z} is quite intuitive, where we simply generate N rules with RNN_{θ} to form \mathbf{z} .

Reasoning Predictor with Logic Rules. The reasoning predictor defines $p_w(\mathbf{a} | \mathcal{G}, \mathbf{q}, \mathbf{z})$. For a query \mathbf{q} , the predictor uses a set of rules \mathbf{z} to reason on a knowledge graph \mathcal{G} and predict the answer \mathbf{a} .

Following stochastic logic programming (Cussens, 2000), a principled reasoning framework, we use a log-linear model for reasoning. For each query $\mathbf{q} = (h, \mathfrak{r}, ?)$, a compositional rule is able to find different grounding paths on graph \mathcal{G} , leading to different candidate answers. For example, given query $(\text{Alice}, \text{hobby}, ?)$, a rule $\text{hobby} \leftarrow \text{friend} \wedge \text{hobby}$ can have two groundings, $\text{hobby}(\text{Alice}, \text{Sing}) \leftarrow \text{friend}(\text{Alice}, \text{Bob}) \wedge \text{hobby}(\text{Bob}, \text{Sing})$ and $\text{hobby}(\text{Alice}, \text{Ski}) \leftarrow \text{friend}(\text{Alice}, \text{Charlie}) \wedge \text{hobby}(\text{Charlie}, \text{Ski})$, yielding two candidate answers *Sing* and *Ski*.

Let \mathcal{A} be the set of candidate answers which can be discovered by any logic rule in the set \mathbf{z} . For each candidate answer $e \in \mathcal{A}$, we compute the following scalar $\text{score}_w(e)$ for that candidate:

$$\text{score}_w(e) = \sum_{\text{rule} \in \mathbf{z}} \text{score}_w(e | \text{rule}) = \sum_{\text{rule} \in \mathbf{z}} \sum_{\text{path} \in \mathcal{P}(h, \text{rule}, e)} \psi_w(\text{rule}) \cdot \phi_w(\text{path}), \quad (4)$$

where $\mathcal{P}(h, \text{rule}, e)$ is the set of grounding paths which start at h and end at e following a *rule* (e.g., $\text{Alice} \xrightarrow{\text{friend}} \text{Bob} \xrightarrow{\text{hobby}} \text{Sing}$). $\psi_w(\text{rule})$ and $\phi_w(\text{path})$ are scalar weights of each *rule* and *path*. Intuitively, the score of each candidate answer e is the sum of scores contributed by each rule, i.e., $\text{score}_w(e | \text{rule})$. To get $\text{score}_w(e | \text{rule})$, we sum over every grounding *path* found in the graph \mathcal{G} .

For the scalar weight $\psi_w(\text{rule})$ of each *rule*, it is a learnable parameter to optimize. For the score $\phi_w(\text{path})$ of each specific *path*, we explore two methods for parameterization. One method always sets $\phi_w(\text{path}) = 1$. However, this method cannot distinguish between different relational paths. To address the problem, for the other method, we follow an embedding algorithm RotatE (Sun et al., 2019) to introduce an embedding for each entity and model each relation as a rotation operator on entity embeddings. Then for each grounding *path* of *rule* starting from h to e , if we rotate the embedding of h according to the rotation operators defined by the body relations of *rule*, we should expect to obtain an embedding close to the embedding of e . Thus we compute the similarity between the derived embedding and the embedding of e as $\phi_w(\text{path})$, which can be viewed as a measure of the soundness and consistency of each *path*. For example, given a path $\text{Alice} \xrightarrow{\text{friend}} \text{Bob} \xrightarrow{\text{hobby}} \text{Sing}$, we rotate *Alice*'s embedding with the operators defined by *friend* and *hobby*. Afterwards we compute the similarity of the derived embedding and embedding of *Sing* as $\phi_w(\text{path})$. Such a method allows us to compute a specific $\phi_w(\text{path})$ for each *path*, which further leads to more precise scores for different candidate answers. See App. C for more details of the parameterization method.

Once we have the score for every candidate answer, we can further define the probability that the answer \mathbf{a} of the query \mathbf{q} is entity e by using a softmax function as follows:

$$p_w(\mathbf{a} = e | \mathcal{G}, \mathbf{q}, \mathbf{z}) = \frac{\exp(\text{score}_w(e))}{\sum_{e' \in \mathcal{A}} \exp(\text{score}_w(e'))}. \quad (5)$$

3.3 OPTIMIZATION

Next, we introduce how we optimize the reasoning predictor and rule generator to maximize the objective in Eq. (2). At each training iteration, we first update the reasoning predictor p_w according to some rules generated by the generator, and then update the rule generator p_θ with an EM algorithm. In the E-step, a set of high-quality rules are identified from all generated rules via posterior inference, with the rule generator as the prior and the reasoning predictor as the likelihood. In the M-step, the rule generator is then updated to be consistent with the high-quality rules selected in the E-step.

Formally, at each training iteration, we start with maximizing the objective $\mathcal{O}(\theta, w)$ in Eq. (2) with respect to the reasoning predictor p_w . To do that, we notice that there is an expectation operation with respect to $p_\theta(\mathbf{z} | \mathbf{q})$ for each training instance $(\mathcal{G}, \mathbf{q}, \mathbf{a})$. By drawing a sample $\hat{\mathbf{z}} \sim p_\theta(\mathbf{z} | \mathbf{q})$ for query \mathbf{q} , we can approximate the objective function of w at each training instance $(\mathcal{G}, \mathbf{q}, \mathbf{a})$ as below:

$$\mathcal{O}_{(\mathcal{G}, \mathbf{q}, \mathbf{a})}(w) = \log \mathbb{E}_{p_\theta(\mathbf{z} | \mathbf{q})} [p_w(\mathbf{a} | \mathcal{G}, \mathbf{q}, \mathbf{z})] \approx \log p_w(\mathbf{a} | \mathcal{G}, \mathbf{q}, \hat{\mathbf{z}}) \quad (6)$$

Basically, we sample a set of rules $\hat{\mathbf{z}}$ from the generator and feed $\hat{\mathbf{z}}$ into the reasoning predictor. Then the parameter w of the reasoning predictor is updated to maximize the log-likelihood of the answer \mathbf{a} .

With the updated reasoning predictor, we then update the rule generator p_θ to maximize the objective $\mathcal{O}(\theta, w)$. In general, this can be achieved by REINFORCE (Williams, 1992) or the reparameterization trick (Jang et al., 2017; Maddison et al., 2017), but they are less effective in our problem due to the large number of logic rules. Therefore, an EM framework is developed to optimize the rule generator.

E-step. Recall that when optimizing the reasoning predictor, we draw a set of rules $\hat{\mathbf{z}}$ for each data instance $(\mathcal{G}, \mathbf{q}, \mathbf{a})$, and let the reasoning predictor use $\hat{\mathbf{z}}$ to predict \mathbf{a} . For each data instance, the E-step aims to identify a set of K high-quality rules \mathbf{z}_I from all generated rules $\hat{\mathbf{z}}$, i.e., $\mathbf{z}_I \subset \hat{\mathbf{z}}, |\mathbf{z}_I| = K$.

Formally, this is achieved by considering the posterior probability of each subset of logic rules \mathbf{z}_I , i.e., $p_{\theta, w}(\mathbf{z}_I | \mathcal{G}, \mathbf{q}, \mathbf{a}) \propto p_w(\mathbf{a} | \mathcal{G}, \mathbf{q}, \mathbf{z}_I) p_\theta(\mathbf{z}_I | \mathcal{G}, \mathbf{q})$, with prior of \mathbf{z}_I from the rule generator p_θ and likelihood from the reasoning predictor p_w . The posterior combines knowledge from both the rule generator and reasoning predictor, so the likely set of high-quality rules can be obtained by sampling from the posterior. However, sampling from the posterior is nontrivial due to its intractable partition function, so we approximate the posterior using a more tractable form with the proposition below:

Proposition 1 Consider a data instance $(\mathcal{G}, \mathbf{q}, \mathbf{a})$ with $\mathbf{q} = (h, r, ?)$ and $\mathbf{a} = t$. For a set of rules $\hat{\mathbf{z}}$ generated by the rule generator p_θ , we can compute the following score H for each rule $e \in \hat{\mathbf{z}}$:

$$H(\text{rule}) = \left\{ \text{score}_w(t | \text{rule}) - \frac{1}{|\mathcal{A}|} \sum_{e \in \mathcal{A}} \text{score}_w(e | \text{rule}) \right\} + \log \text{RNN}_\theta(\text{rule} | r), \quad (7)$$

where \mathcal{A} is the set of all candidate answers discovered by rules in $\hat{\mathbf{z}}$, $\text{score}_w(e | \text{rule})$ is the score that each rule contributes to entity e as defined by Eq. (4), $\text{RNN}_\theta(\text{rule} | r)$ is the prior probability of rule computed by the generator. Suppose $s = \max_{e \in \mathcal{A}} |\text{score}_w(e)| < 1$. Then for a subset of rules $\mathbf{z}_I \subset \hat{\mathbf{z}}$ with $|\mathbf{z}_I| = K$, the log-probability $\log p_{\theta, w}(\mathbf{z}_I | \mathcal{G}, \mathbf{q}, \mathbf{a})$ can be approximated as follows:

$$\left| \log p_{\theta, w}(\mathbf{z}_I | \mathcal{G}, \mathbf{q}, \mathbf{a}) - \left(\sum_{\text{rule} \in \mathbf{z}_I} H(\text{rule}) + \gamma(\mathbf{z}_I) + \text{const} \right) \right| \leq s^2 + O(s^4) \quad (8)$$

where const is a constant term that is independent from \mathbf{z}_I , $\gamma(\mathbf{z}_I) = \log(K! / \prod_{\text{rule} \in \mathbf{z}_I} n_{\text{rule}}!)$, with K being the given size of set \mathbf{z}_I and n_{rule} being the number of times each rule appears in \mathbf{z}_I .

We prove the proposition in App. A.1. In practice, we can apply weight decay to the weight of logic rules in Eq. (4), and thereby reduce $s = \max_{e \in \mathcal{A}} |\text{score}_w(e)|$ to get a more precise approximation.

The above proposition allows us to utilize $(\sum_{\text{rule} \in \mathbf{z}_I} H(\text{rule}) + \gamma(\mathbf{z}_I) + \text{const})$ to approximate the log-posterior $\log p_{\theta, w}(\mathbf{z}_I | \mathcal{G}, \mathbf{q}, \mathbf{a})$, yielding a distribution $q(\mathbf{z}_I) \propto \exp(\sum_{\text{rule} \in \mathbf{z}_I} H(\text{rule}) + \gamma(\mathbf{z}_I))$ as a

good approximation of the posterior. It turns out that the derived $q(z_I)$ is a multinomial distribution, and thus sampling from $q(z_I)$ is more tractable. Specifically, a sample \hat{z}_I from $q(z_I)$ can be formed with K logic rules which are independently sampled from \hat{z} , where the probability of sampling each $rule$ is computed as $\exp(H(rule))/(\sum_{rule' \in \hat{z}} \exp(H(rule')))$. We provide the proof in the App. A.2.

Intuitively, we could view $H(rule)$ of each $rule$ as an assessment of the rule quality, which considers two factors. The first factor is based on the reasoning predictor p_w , and it is computed as the score that a $rule$ contributes to the correct answer t minus the mean score that this $rule$ contributes to other candidate answers. If a rule gives higher score to the true answer and lower score to other candidate answers, then the rule is more likely to be important. The second factor is based on the rule generator p_θ , where we compute the prior probability for each $rule$ and use the probability for regularization.

Empirically, we find that picking K rules with highest $H(rule)$ to form \hat{z}_I works better than sampling from the posterior. Similar observations have been made on the node classification task (Qu et al., 2019). In fact, \hat{z}_I formed by the top- K rules is an MAP estimation of the posterior, and thus the variant of picking top- K rules yields a hard-assignment EM algorithm (Koller & Friedman, 2009). Despite the reduced theoretical guarantees, we use this variant in practice for its good performance.

M-step. Once we obtain a set of high-quality logic rules \hat{z}_I for each data instance $(\mathcal{G}, \mathbf{q}, \mathbf{a})$ in the E-step, we further leverage those rules to update the parameters θ of the rule generator in the M-step.

Specifically, for each data instance $(\mathcal{G}, \mathbf{q}, \mathbf{a})$, we treat the corresponding rule set \hat{z}_I as part of the (now complete) training data, and update the rule generator by maximizing the log-likelihood of \hat{z}_I :

$$\mathcal{O}_{(\mathcal{G}, \mathbf{q}, \mathbf{a})}(\theta) = \log p_\theta(\hat{z}_I | \mathbf{q}) = \sum_{rule \in \hat{z}_I} \log \text{RNN}_\theta(rule | r) + \text{const.} \quad (9)$$

With the above objective, the feedback from the reasoning predictor can be effectively distilled into the rule generator. In this way, the rule generator will learn to only generate high-quality rules for the reasoning predictor to explore, which reduces the search space and yields better empirical results.

For more detailed analysis of the EM algorithm for optimizing $\mathcal{O}(\theta, w)$, please refer to App. B.

Algorithm 1 Workflow of RNNLogic

while not converge **do**

 For each instance, use the rule generator p_θ to generate a set of rules \hat{z} ($|\hat{z}| = N$).

 For each instance, update the reasoning predictor p_w based on generated rules \hat{z} and Eq. (6).

 □ *E-step*:

 For each instance, identify K high-quality rules \hat{z}_I from \hat{z} according to $H(rule)$ in Eq. (7).

 □ *M-step*:

 For each instance, update the rule generator p_θ according to the identified rules and Eq. (9).

end while

During testing, for each query, use p_θ to generate N rules and feed them into p_w for prediction.

4 EXPERIMENT

4.1 EXPERIMENT SETTINGS

Datasets. We choose four datasets for evaluation, including FB15k-237 (Toutanova & Chen, 2015), WN18RR (Dettmers et al., 2018), Kinship and UMLS (Kok & Domingos, 2007). For Kinship and UMLS, there are no standard data splits, so we randomly sample 30% of all the triplets for training, 20% for validation, and the rest 50% for testing. The detailed statistics are summarized in the App. D.

Compared Algorithms. We compare the following algorithms in experiment:

□ *Rule learning methods.* For traditional statistical relational learning methods, we choose Markov logic networks (Richardson & Domingos, 2006), boosted relational dependency networks (Natarajan et al., 2010) and path ranking (Lao & Cohen, 2010). We also consider neural logic programming methods, including NeuralLP (Yang et al., 2017), DRUM (Sadeghian et al., 2019) and NLIL (Yang & Song, 2020). In addition, we compare against CTP (Minervini et al., 2020), a differentiable method based on neural theorem provers. Besides, we consider three reinforcement learning methods, which

Table 1: Results of reasoning on FB15k-237 and WN18RR. $H@k$ is in %. [*] means the numbers are taken from the original papers. [†] means we rerun the methods with the same evaluation process.

Category	Algorithm	FB15k-237					WN18RR				
		MR	MRR	H@1	H@3	H@10	MR	MRR	H@1	H@3	H@10
No Rule Learning	TransE*	357	0.294	-	-	46.5	3384	0.226	-	-	50.1
	DistMult*	254	0.241	15.5	26.3	41.9	5110	0.43	39	44	49
	ComplEx*	339	0.247	15.8	27.5	42.8	5261	0.44	41	46	51
	ComplEx-N3*	-	0.37	-	-	56	-	0.48	-	-	57
	ConvE*	244	0.325	23.7	35.6	50.1	4187	0.43	40	44	52
	TuckER*	-	0.358	26.6	39.4	54.4	-	0.470	44.3	48.2	52.6
	RotatE*	177	0.338	24.1	37.5	53.3	3340	0.476	42.8	49.2	57.1
Rule Learning	PathRank	-	0.087	7.4	9.2	11.2	-	0.189	17.1	20.0	22.5
	NeuralLP†	-	0.237	17.3	25.9	36.1	-	0.381	36.8	38.6	40.8
	DRUM†	-	0.238	17.4	26.1	36.4	-	0.382	36.9	38.8	41.0
	NLIL*	-	0.25	-	-	32.4	-	-	-	-	-
	M-Walk*	-	0.232	16.5	24.3	-	-	0.437	41.4	44.5	-
RNNLogic	w/o emb.	538	0.288	20.8	31.5	44.5	7527	0.455	41.4	47.5	53.1
	with emb.	232	0.344	25.2	38.0	53.0	4615	0.483	44.6	49.7	55.8

Table 2: Results of knowledge graph reasoning on the FB15k-237 and WN18RR datasets with only $(h, r, ?)$ -queries. $H@k$ is in %. [*] means that the numbers are taken from the original papers.

Category	Algorithm	FB15k-237					WN18RR				
		MR	MRR	H@1	H@3	H@10	MR	MRR	H@1	H@3	H@10
Rule Learning	MINERVA*	-	0.293	21.7	32.9	45.6	-	0.448	41.3	45.6	51.3
	MultiHopKG*	-	0.407	32.7	-	56.4	-	0.472	43.7	-	54.2
RNNLogic	w/o emb.	459.0	0.377	28.9	41.2	54.9	7662.8	0.478	43.8	50.3	55.3
	with emb.	146.1	0.443	34.4	48.9	64.0	3767.0	0.506	46.3	52.3	59.2

are MINERVA (Das et al., 2018), MultiHopKG (Lin et al., 2018) and M-Walk (Shen et al., 2018).

□ *Other methods.* We also compare with some embedding methods, including TransE (Bordes et al., 2013), DistMult (Yang et al., 2015), ComplEx (Trouillon et al., 2016), ComplEx-N3 (Lacroix et al., 2018), ConvE (Dettmers et al., 2018), TuckER (Balazevic et al., 2019) and RotatE (Sun et al., 2019). □ *RNNLogic.* For RNNLogic, we consider two model variants. The first variant assigns a constant score to different grounding paths in the reasoning predictor, i.e., $\phi_w(path) = 1$ in Eq. (4), and we denote this variant as *w/o emb.*. The second variant leverages entity embeddings and relation embeddings to compute the path score $\phi_w(path)$, and we denote the variant as *with emb.*

Evaluation Metrics. During evaluation, for each test triplet (h, r, t) , we build two queries $(h, r, ?)$ and $(t, r^{-1}, ?)$ with answers t and h . For each query, we compute a probability for each entity, and compute the rank of the correct answer. Given the ranks from all queries, we report the Mean Rank (MR), Mean Reciprocal Rank (MRR) and Hit@ k ($H@k$) under the filtered setting (Bordes et al., 2013), which is used by most existing studies. Note that there can be a case where an algorithm assigns the same probability to the correct answer and a few other entities. For such a case, many methods compute the rank of the correct answer as $(m + 1)$ where m is the number of entities receiving higher probabilities than the correct answer. This setup can be problematic according to Sun et al. (2020). For fair comparison, in that case we compute the expectation of each evaluation metric over all the random shuffles of entities which receive the same probability as the correct answer. For example, if there are n entities which have the same probability as the correct answer in the above case, then we treat the rank of the correct answer as $(m + (n + 1)/2)$ when computing Mean Rank.

Besides, we notice that in MINERVA (Das et al., 2018) and MultiHopKG (Lin et al., 2018), they only consider queries in the form of $(h, r, ?)$, which is different from our default setting. To make fair comparison with these methods, we also apply RNNLogic to this setting and report the performance.

Experimental Setup of RNNLogic. For each training triplet (h, r, t) , we add an inverse triplet (t, r^{-1}, h) into the training set, yielding an augmented set of training triplets \mathcal{T} . We use a closed-world assumption for model training, which assumes that any triplet outside \mathcal{T} is incorrect. To build a training instance from p_{data} , we first randomly sample a triplet (h, r, t) from \mathcal{T} , and then form an instance as $(\mathcal{G} = \mathcal{T} \setminus \{(h, r, t)\}, \mathbf{q} = (h, r, ?), \mathbf{a} = t)$. Basically, we use the sampled triplet (h, r, t) to construct the query and answer, and use the rest of triplets in \mathcal{T} to form the background knowledge graph \mathcal{G} . During testing, the background knowledge graph \mathcal{G} is formed with all the triplets in \mathcal{T} . For the rule generator, the maximum length of generated rules is set to 4 for FB15k-237, 5 for WN18RR, and 3 for the rest, which are selected on validation data. See App. D for the detailed setting.

Table 3: Results of reasoning on the Kinship and UMLS datasets. H@k is in %.

Category	Algorithm	Kinship					UMLS				
		MR	MRR	H@1	H@3	H@10	MR	MRR	H@1	H@3	H@10
No Rule Learning	DistMult	8.5	0.354	18.9	40.0	75.5	14.6	0.391	25.6	44.5	66.9
	ComplEx	7.8	0.418	24.2	49.9	81.2	13.6	0.411	27.3	46.8	70.0
	ComplEx-N3	-	0.605	43.7	71.0	92.1	-	0.791	68.9	87.3	95.7
	TuckER	6.2	0.603	46.2	69.8	86.3	5.7	0.732	62.5	81.2	90.9
	RotatE	3.7	0.651	50.4	75.5	93.2	4.0	0.744	63.6	82.2	93.9
Rule Learning	MLN	10.0	0.351	18.9	40.8	70.7	7.6	0.688	58.7	75.5	86.9
	Boosted RDN	25.2	0.469	39.5	52.0	56.7	54.8	0.227	14.7	25.6	37.6
	PathRank	-	0.369	27.2	41.6	67.3	-	0.197	14.8	21.4	25.2
	NeuralLP	16.9	0.302	16.7	33.9	59.6	10.3	0.483	33.2	56.3	77.5
	DRUM	11.6	0.334	18.3	37.8	67.5	8.4	0.548	35.8	69.9	85.4
	MINERVA	-	0.401	23.5	46.7	76.6	-	0.564	42.6	65.8	81.4
CTP	-	0.335	17.7	37.6	70.3	-	0.404	28.8	43.0	67.4	
RNNLogic	w/o emb.	3.9	0.639	49.5	73.1	92.4	5.3	0.745	63.0	83.3	92.4
	with emb.	3.1	0.722	59.8	81.4	94.9	3.1	0.842	77.2	89.1	96.5

4.2 RESULTS

Comparison against Existing Methods. We present the results on the FB15k-237 and WN18RR datasets in Tab. 1 and Tab. 2. The results on the Kinship and UMLS datasets are shown in Tab. 3.

We first compare RNNLogic with rule learning methods. RNNLogic achieves much better results than statistical relational learning methods (MLN, Boosted RDN, PathRank) and neural differentiable methods (NeuralLP, DRUM, NLIL, CTP). This is because the rule generator and reasoning predictor of RNNLogic can collaborate with each other to reduce search space and learn better rules. RNNLogic also outperforms reinforcement learning methods (MINERVA, MultiHopKG, M-Walk). The reason is that RNNLogic is optimized with an EM-based framework, in which the reasoning predictor provides more useful feedback to the rule generator, and thus addresses the challenge of sparse reward.

We then compare RNNLogic against state-of-the-art embedding-based methods. For RNNLogic with embeddings in the reasoning predictor (*with emb.*), it outperforms most compared methods in most cases, and the reason is that RNNLogic is able to use logic rules to enhance reasoning performance. For RNNLogic without embedding (*w/o emb.*), it achieves comparable results to embedding-based methods, especially on WN18RR, Kinship and UMLS where the training triplets are quite limited.

Quality of Learned Logic Rules. Next, we study the quality of rules learned by different methods for reasoning. For each trained model, we let it generate I rules with highest qualities per query relation, and use them to train a predictor *w/o emb.* as in Eq. (5) for reasoning. For RNNLogic, the quality of each rule is measured by its prior probability from the rule generator, and we use beam search to infer top- I rules. The results at different I are in Fig. 2, where RNNLogic achieves much better results. Even with only 10 rules per relation, RNNLogic still achieves competitive results.

Performance w.r.t. the Number of Training Triplets. To better evaluate different methods under cases where training triplets are very limited, in this section we reduce the amount of training data on Kinship and UMLS to see how the performance varies. The results are presented in Fig. 4. We see that RNNLogic *w/o emb.* achieves the best results. Besides, the improvement over RotatE is more significant as we reduce training triplets, showing that RNNLogic is more robust to data sparsity.







Figure 2: Performance w.r.t. # logic rules. RNNLogic achieves competitive results even with 10 rules per query relation.

Figure 3: Performance w.r.t. embedding dimension.

Performance w.r.t. Embedding Dimension. RNNLogic *with emb.* uses entity and relation embeddings to improve the reasoning predictor. Next, we study its performance with different embedding dimensions. The results are presented in Fig. 3, where we compare against RotatE (Sun et al., 2019). We see that RNNLogic significantly outperforms RotatE at every embedding dimension. The improvement is mainly from the use of logic rules, showing that our learned rules are indeed helpful.

Comparison of Optimization Algorithms. RNNLogic uses an EM algorithm to optimize the rule generator. In practice, the generator can also be optimized with REINFORCE (Williams, 1992) (see App. F for details). We empirically compare the two algorithms in the *w/o emb.* case. The results on Kinship and UMLS are presented in Tab. 4. We see EM consistently outperforms REINFORCE.



	Kinship MRR	UMLS MRR
REINFORCE	0.312	0.504
EM	0.639	0.745

Table 4: Comparison between REINFORCE and EM.

Figure 4: Performance w.r.t. # training triplets. RNNLogic is more robust to data sparsity even without using embeddings.

Case Study of Generated Logic Rules. Finally, we show some logic rules generated by RNNLogic on the FB15k-237 dataset in Tab. 5. We can see that these logic rules are meaningful and diverse. The first rule is a subrelation rule. The third and fifth rules are two-hop compositional rules. The rest of logic rules have even more complicated forms. This case study shows that RNNLogic can indeed learn useful and diverse rules for reasoning. For more generated logic rules, please refer to App. E.

Table 5: Case study of the rules generated by the rule generator.

$Appears_in_TV_Show(X, Y) \leftarrow Actor_of(X, Y)$
$Appears_in_TV_Show(X, Y) \leftarrow Creator_of(X, U) \wedge Has_Producer(U, V) \wedge Appears_in_TV_Show(V, Y)$
$ORG_in_State(X, Y) \leftarrow ORG_in_City(X, U) \wedge City_Locates_in_State(U, Y)$
$ORG_in_State(X, Y) \leftarrow ORG_in_City(X, U) \wedge Address_of_PERS.(U, V) \wedge Born_in(V, W) \wedge Town_in_State(W, Y)$
$Person_Nationality(X, Y) \leftarrow Born_in(X, U) \wedge Place_in_Country(U, Y)$
$Person_Nationality(X, Y) \leftarrow Student_of_Educational_Institution(X, U) \wedge ORG_Endowment_Currency(U, V) \wedge Currency_Used_in_Region(V, W) \wedge Region_in_Country(W, Y)$

5 CONCLUSION

This paper studies learning logic rules for knowledge graph reasoning, and an approach called RNNLogic is proposed. RNNLogic treats a set of logic rules as a latent variable, and a rule generator as well as a reasoning predictor with logic rules are jointly learned. We develop an EM-based algorithm for optimization. Extensive experiments prove the effectiveness of RNNLogic. In the future, we plan to study generating more complicated logic rules rather than only compositional rules. Besides, we plan to extend RNNLogic to other reasoning problems, such as question answering.

ACKNOWLEDGMENTS

This project is supported by the Natural Sciences and Engineering Research Council (NSERC) Discovery Grant, the Canada CIFAR AI Chair Program, collaboration grants between Microsoft Research and Mila, Samsung Electronics Co., Ltd., Amazon Faculty Research Award, Tencent AI Lab Rhino-Bird Gift Fund and a NRC Collaborative R&D Project (AI4D-CORE-06). This project was also partially funded by IVADO Fundamental Research Project grant PRF-2019-3583139727.

REFERENCES

- Ivana Balazevic, Carl Allen, and Timothy Hospedales. Tucker: Tensor factorization for knowledge graph completion. In *EMNLP*, 2019.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. Translating embeddings for modeling multi-relational data. In *NeurIPS*, 2013.
- Liwei Cai and William Yang Wang. Kbgan: Adversarial learning for knowledge graph embeddings. In *NAACL*, 2018.
- Wenhu Chen, Wenhan Xiong, Xifeng Yan, and William Wang. Variational knowledge graph reasoning. In *NAACL*, 2018.
- William W Cohen, Fan Yang, and Kathryn Rivard Mazaitis. Tensorlog: Deep learning meets probabilistic databases. *Journal of Artificial Intelligence Research*, 2018.
- James Cussens. Stochastic logic programs: sampling, inference and applications. In *UAI*, 2000.
- Rajarshi Das, Shehzaad Dhuliawala, Manzil Zaheer, Luke Vilnis, Ishan Durugkar, Akshay Krishnamurthy, Alex Smola, and Andrew McCallum. Go for a walk and arrive at the answer: Reasoning over paths in knowledge bases using reinforcement learning. In *ICLR*, 2018.
- Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. Convolutional 2d knowledge graph embeddings. In *AAAI*, 2018.
- Jonathan Eckstein, Noam Goldberg, and Ai Kagawa. Rule-enhanced penalized regression by column generation using rectangular maximum agreement. In *ICML*, 2017.
- Luis Antonio Galárraga, Christina Teflioudi, Katja Hose, and Fabian Suchanek. Amie: association rule mining under incomplete evidence in ontological knowledge bases. In *WWW*, 2013.
- Noam Goldberg and Jonathan Eckstein. Boosting classifiers with tightened l0-relaxation penalties. In *ICML*, 2010.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 1997.
- Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *ICLR*, 2017.
- Tushar Khot, Sriraam Natarajan, Kristian Kersting, and Jude Shavlik. Learning markov logic networks via functional gradient boosting. In *ICDM*, 2011.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2014.
- Stanley Kok and Pedro Domingos. Learning the structure of markov logic networks. In *ICML*, 2005.
- Stanley Kok and Pedro Domingos. Statistical predicate invention. In *ICML*, 2007.
- Daphne Koller and Nir Friedman. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- Timothee Lacroix, Nicolas Usunier, and Guillaume Obozinski. Canonical tensor decomposition for knowledge base completion. In *ICML*, 2018.
- Ni Lao and William W Cohen. Relational retrieval using a combination of path-constrained random walks. *Machine learning*, 2010.
- Ni Lao, Tom Mitchell, and William W Cohen. Random walk inference and learning in a large scale knowledge base. In *EMNLP*, 2011.
- Xi Victoria Lin, Richard Socher, and Caiming Xiong. Multi-hop knowledge graph reasoning with reward shaping. In *EMNLP*, 2018.
- Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. In *ICLR*, 2017.

- Christian Meilicke, Melisachew Wudage Chekol, Daniel Ruffinelli, and Heiner Stuckenschmidt. Anytime bottom-up rule learning for knowledge graph completion. In *IJCAI*, 2019.
- Pasquale Minervini, Sebastian Riedel, Pontus Stenetorp, Edward Grefenstette, and Tim Rocktäschel. Learning reasoning strategies in end-to-end differentiable proving. In *ICML*, 2020.
- Sriraam Natarajan, Tushar Khot, Kristian Kersting, Bernd Gutmann, and Jude Shavlik. Boosting relational dependency networks. In *ICILP*, 2010.
- Radford M Neal and Geoffrey E Hinton. A view of the em algorithm that justifies incremental, sparse, and other variants. In *Learning in graphical models*. Springer, 1998.
- Jennifer Neville and David Jensen. Relational dependency networks. *Journal of Machine Learning Research*, 2007.
- Maximilian Nickel, Lorenzo Rosasco, and Tomaso Poggio. Holographic embeddings of knowledge graphs. In *AAAI*, 2016.
- Meng Qu and Jian Tang. Probabilistic logic neural networks for reasoning. In *NeurIPS*, 2019.
- Meng Qu, Yoshua Bengio, and Jian Tang. Gmnn: Graph markov neural networks. In *ICML*, 2019.
- Matthew Richardson and Pedro Domingos. Markov logic networks. *Machine learning*, 2006.
- Tim Rocktäschel and Sebastian Riedel. End-to-end differentiable proving. In *NeurIPS*, 2017.
- Ali Sadeghian, Mohammadreza Armandpour, Patrick Ding, and Daisy Zhe Wang. Drum: End-to-end differentiable rule mining on knowledge graphs. In *NeurIPS*, 2019.
- Yelong Shen, Jianshu Chen, Po-Sen Huang, Yuqing Guo, and Jianfeng Gao. M-walk: Learning to walk over graphs using monte carlo tree search. In *NeurIPS*, 2018.
- Slavko Simic. On a global upper bound for jensen’s inequality. *Journal of mathematical analysis and applications*, 2008.
- Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. Rotate: Knowledge graph embedding by relational rotation in complex space. In *ICLR*, 2019.
- Zhiqing Sun, Shikhar Vashishth, Soumya Sanyal, Partha Talukdar, and Yiming Yang. A re-evaluation of knowledge graph completion methods. In *ACL*, 2020.
- Komal K Teru and William L Hamilton. Inductive relation prediction on knowledge graphs. In *ICML*, 2020.
- Kristina Toutanova and Danqi Chen. Observed versus latent features for knowledge base and text inference. In *Workshop on Continuous Vector Space Models and their Compositionality*, 2015.
- Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. Complex embeddings for simple link prediction. In *ICML*, 2016.
- William Yang Wang, Kathryn Mazaitis, and William W Cohen. Programming with personalized pagerank: a locally groundable first-order probabilistic logic. In *CIKM*, 2013.
- William Yang Wang, Kathryn Mazaitis, and William W Cohen. Proppr: Efficient first-order probabilistic logic programming for structure discovery, parameter learning, and scalable inference. In *Workshops at AAAI*, 2014a.
- William Yang Wang, Kathryn Mazaitis, and William W Cohen. Structure learning via parameter learning. In *CIKM*, 2014b.
- Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. Knowledge graph embedding by translating on hyperplanes. In *AAAI*, 2014c.
- Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 1992.

Wenhan Xiong, Thien Hoang, and William Yang Wang. Deeppath: A reinforcement learning method for knowledge graph reasoning. In *EMNLP*, 2017.

Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding entities and relations for learning and inference in knowledge bases. In *ICLR*, 2015.

Fan Yang, Zhilin Yang, and William W Cohen. Differentiable learning of logical rules for knowledge base reasoning. In *NeurIPS*, 2017.

Yuan Yang and Le Song. Learn to explain efficiently via neural logic inductive learning. In *ICLR*, 2020.

Yuyu Zhang, Xinshi Chen, Yuan Yang, Arun Ramamurthy, Bo Li, Yuan Qi, and Le Song. Efficient probabilistic logic reasoning with graph neural networks. In *ICLR*, 2020.

A PROOFS

This section presents the proofs of some propositions used in the optimization algorithm of RNNLogic.

Recall that in the E-step of the optimization algorithm, we aim to sample from the posterior distribution over rule sets. However, directly sampling from the posterior distribution is intractable due to the intractable partition function. Therefore, we introduce Proposition 1, which gives an approximation distribution with more tractable form for the posterior distribution. With this approximation, sampling becomes much easier. In Section A.1, we present the proof of Proposition 1. In Section A.2, we show how to perform sampling based on the approximation of the posterior distribution.

A.1 PROOF OF PROPOSITION 1

Next, we prove proposition 1, which is used to approximate the true posterior probability in the E-step of optimization. We first restate the proposition as follows:

Proposition Consider a data instance $(\mathcal{G}, \mathbf{q}, \mathbf{a})$ with $\mathbf{q} = (h, r, ?)$ and $\mathbf{a} = t$. For a set of rules $\hat{\mathbf{z}}$ generated by the rule generator p_θ , we can compute the following score H for each rule $e \in \hat{\mathbf{z}}$:

$$H(\text{rule}) = \left\{ \text{score}_w(t|\text{rule}) - \frac{1}{|\mathcal{A}|} \sum_{e \in \mathcal{A}} \text{score}_w(e|\text{rule}) \right\} + \log \text{RNN}_\theta(\text{rule}|r),$$

where \mathcal{A} is the set of all candidate answers discovered by rules in $\hat{\mathbf{z}}$, $\text{score}_w(e|\text{rule})$ is the score that each rule contributes to entity e as defined by Eq. (4), $\text{RNN}_\theta(\text{rule}|r)$ is the prior probability of rule computed by the generator. Suppose $s = \max_{e \in \mathcal{A}} |\text{score}_w(e)| < 1$. Then for a subset of rules $\mathbf{z}_I \subset \hat{\mathbf{z}}$ with $|\mathbf{z}_I| = K$, the log-probability $\log p_{\theta,w}(\mathbf{z}_I|\mathcal{G}, \mathbf{q}, \mathbf{a})$ could be approximated as follows:

$$\left| \log p_{\theta,w}(\mathbf{z}_I|\mathcal{G}, \mathbf{q}, \mathbf{a}) - \left(\sum_{\text{rule} \in \mathbf{z}_I} H(\text{rule}) + \gamma(\mathbf{z}_I) + \text{const} \right) \right| \leq s^2 + O(s^4)$$

where const is a constant term that is independent from \mathbf{z}_I , $\gamma(\mathbf{z}_I) = \log(K! / \prod_{\text{rule} \in \hat{\mathbf{z}}} n_{\text{rule}}!)$, with K being the given size of set \mathbf{z}_I and n_{rule} being the number of times each rule appears in \mathbf{z}_I .

Proof: We first rewrite the posterior probability as follows:

$$\begin{aligned} \log p_{\theta,w}(\mathbf{z}_I|\mathcal{G}, \mathbf{q}, \mathbf{a}) &= \log p_w(\mathbf{a}|\mathcal{G}, \mathbf{q}, \mathbf{z}_I) + \log p_\theta(\mathbf{z}_I|\mathbf{q}) + \text{const} \\ &= \log \frac{\exp(\text{score}_w(t))}{\sum_{e \in \mathcal{A}} \exp(\text{score}_w(e))} + \log \text{Mu}(\mathbf{z}_I|K, \text{RNN}_\theta(\cdot|r)) + \text{const}, \end{aligned}$$

where const is a constant term which does not depend on the choice of \mathbf{z}_I , and $\text{RNN}_\theta(\cdot|r)$ defines a probability distribution over all the composition-based logic rules. The probability mass function of the above multinomial distribution $\text{Mu}(\mathbf{z}_I|K, \text{RNN}_\theta(\cdot|r))$ can be written as below:

$$\text{Mu}(\mathbf{z}_I|K, \mathbf{q}) = \frac{K!}{\prod_{\text{rule} \in \hat{\mathbf{z}}} n_{\text{rule}}!} \prod_{\text{rule} \in \hat{\mathbf{z}}} \text{RNN}_\theta(\text{rule}|r)^{n_{\text{rule}}},$$

where K is the size of set \mathbf{z}_I and n_{rule} is the number of times a rule appears in \mathbf{z}_I .

Letting $\gamma(\mathbf{z}_I) = \log(K! / \prod_{\text{rule} \in \hat{\mathbf{z}}} n_{\text{rule}}!)$, then the posterior probability can then be rewritten as:

$$\begin{aligned} &\log p_{\theta,w}(\mathbf{z}_I|\mathcal{G}, \mathbf{q}, \mathbf{a}) \\ &= \log \frac{\exp(\text{score}_w(t))}{\sum_{e \in \mathcal{A}} \exp(\text{score}_w(e))} + \log \frac{K!}{\prod_{\text{rule} \in \hat{\mathbf{z}}} n_{\text{rule}}!} + \log \prod_{\text{rule} \in \hat{\mathbf{z}}} \text{RNN}_\theta(\text{rule}|r)^{n_{\text{rule}}} + \text{const} \\ &= \log \frac{\exp(\text{score}_w(t))}{\sum_{e \in \mathcal{A}} \exp(\text{score}_w(e))} + \gamma(\mathbf{z}_I) + \sum_{\text{rule} \in \mathbf{z}_I} \log \text{RNN}_\theta(\text{rule}|r) + \text{const} \\ &= \text{score}_w(t) - \log \sum_{e \in \mathcal{A}} \exp(\text{score}_w(e)) + \gamma(\mathbf{z}_I) + \sum_{\text{rule} \in \mathbf{z}_I} \log \text{RNN}_\theta(\text{rule}|r) + \text{const}. \end{aligned}$$

The above term $\log \sum_{e \in \mathcal{A}} \exp(\text{score}_w(e))$ makes the posterior distribution hard to deal with, and thus we approximate it using Lemma 1, which we prove at the end of this section.

Lemma 1 Let $e \in \mathcal{A}$ be a finite set of entities, let $|\text{score}_w(e)| \leq s < 1$, and let score_w be a function from entities to real numbers. Then the following inequalities hold:

$$0 \leq \log \left(\sum_{e \in \mathcal{A}} \exp(\text{score}_w(e)) \right) - \left(\sum_{e \in \mathcal{A}} \frac{1}{|\mathcal{A}|} \text{score}_w(e) + \log(|\mathcal{A}|) \right) \leq s^2 + O(s^4).$$

Hence, using the lemma we can get the following upper bound of the posterior probability:

$$\begin{aligned} & \log p_{\theta,w}(\mathbf{z}_I | \mathcal{G}, \mathbf{q}, \mathbf{a}) \\ &= \text{score}_w(t) - \log \sum_{e \in \mathcal{A}} \exp(\text{score}_w(e)) + \gamma(\mathbf{z}_I) + \sum_{rule \in \mathbf{z}_I} \log \text{RNN}_{\theta}(rule|r) + \text{const} \\ &\leq \text{score}_w(t) - \sum_{e \in \mathcal{A}} \frac{1}{|\mathcal{A}|} \text{score}_w(e) + \gamma(\mathbf{z}_I) + \sum_{rule \in \mathbf{z}_I} \log \text{RNN}_{\theta}(rule|r) + \text{const} \\ &= \sum_{rule \in \mathbf{z}_I} H(rule) + \gamma(\mathbf{z}_I) + \text{const}, \end{aligned}$$

and also the following lower bound of the posterior probability:

$$\begin{aligned} & \log p_{\theta,w}(\mathbf{z}_I | \mathcal{G}, \mathbf{q}, \mathbf{a}) \\ &= \text{score}_w(t) - \log \sum_{e \in \mathcal{A}} \exp(\text{score}_w(e)) + \gamma(\mathbf{z}_I) + \sum_{rule \in \mathbf{z}_I} \log \text{RNN}_{\theta}(rule|r) + \text{const} \\ &\geq \text{score}_w(t) - \sum_{e \in \mathcal{A}} \frac{1}{|\mathcal{A}|} \text{score}_w(e) + \gamma(\mathbf{z}_I) + \sum_{rule \in \mathbf{z}_I} \log \text{RNN}_{\theta}(rule|r) + \text{const} - s^2 - O(s^4) \\ &= \sum_{rule \in \mathbf{z}_I} H(rule) + \gamma(\mathbf{z}_I) + \text{const} - s^2 - O(s^4), \end{aligned}$$

where const is a constant term which does not depend on \mathbf{z}_I .

By combining the lower and the upper bound, we get:

$$\left| \log p_{\theta,w}(\mathbf{z}_I | \mathcal{G}, \mathbf{q}, \mathbf{a}) - \left(\sum_{rule \in \mathbf{z}_I} H(rule) + \gamma(\mathbf{z}_I) + \text{const} \right) \right| \leq s^2 + O(s^4)$$

Thus, it only remains to prove Lemma 1 to complete the proof. We use Theorem 1 from (Simic, 2008) as a starting point:

Theorem 1 Suppose that $\tilde{x} = \{x_i\}_{i=1}^n$ represents a finite sequence of real numbers belonging to a fixed closed interval $I = [a, b]$, $a < b$. If f is a convex function on I , then we have that:

$$\frac{1}{n} \sum_{i=1}^n f(x_i) - f\left(\frac{1}{n} \sum_{i=1}^n x_i\right) \leq f(a) + f(b) - 2f\left(\frac{a+b}{2}\right).$$

As $(-\log)$ is convex and $\exp(\text{score}_w(e)) \in [\exp(-s), \exp(s)]$, Theorem 1 gives us that:

$$\begin{aligned} & -\frac{1}{|\mathcal{A}|} \sum_{e \in \mathcal{A}} \log(\exp(\text{score}_w(e))) + \log\left(\frac{1}{|\mathcal{A}|} \sum_{e \in \mathcal{A}} \exp(\text{score}_w(e))\right) \\ &\leq -\log(\exp(-s)) - \log(\exp(s)) + 2\log\left(\frac{\exp(-s) + \exp(s)}{2}\right). \end{aligned}$$

After some simplification, we get:

$$\begin{aligned}
& \log \left(\sum_{e \in \mathcal{A}} \exp(\text{score}_w(e)) \right) \\
& \leq \sum_{e \in \mathcal{A}} \frac{1}{|\mathcal{A}|} \text{score}_w(e) + \log(|\mathcal{A}|) + 2 \log \left(\frac{\exp(-s) + \exp(s)}{2} \right) \\
& = \sum_{e \in \mathcal{A}} \frac{1}{|\mathcal{A}|} \text{score}_w(e) + \log(|\mathcal{A}|) + 2s - 2 \log 2 + 2 \log(1 + \exp(-2s)) \\
& \leq \sum_{e \in \mathcal{A}} \frac{1}{|\mathcal{A}|} \text{score}_w(e) + \log(|\mathcal{A}|) + s^2 + O(s^4),
\end{aligned} \tag{10}$$

where the last inequality is based on Taylor's series $\log(1 + e^x) = \log 2 + \frac{1}{2}x + \frac{1}{8}x^2 + O(x^4)$ with $|x| < 1$. On the other hand, according to the well-known Jensen's inequality, we have:

$$\log \left(\frac{1}{|\mathcal{A}|} \sum_{e \in \mathcal{A}} \exp(\text{score}_w(e)) \right) \geq \frac{1}{|\mathcal{A}|} \sum_{e \in \mathcal{A}} \log(\exp(\text{score}_w(e))),$$

which implies:

$$\log \left(\sum_{e \in \mathcal{A}} \exp(\text{score}_w(e)) \right) \geq \sum_{e \in \mathcal{A}} \frac{1}{|\mathcal{A}|} \text{score}_w(e) + \log(|\mathcal{A}|). \tag{11}$$

By combining Eq. (10) and Eq. (11), we obtain:

$$0 \leq \log \left(\sum_{e \in \mathcal{A}} \exp(\text{score}_w(e)) \right) - \left(\sum_{e \in \mathcal{A}} \frac{1}{|\mathcal{A}|} \text{score}_w(e) + \log(|\mathcal{A}|) \right) \leq s^2 + O(s^4).$$

This completes the proof. \square

A.2 SAMPLING BASED ON THE APPROXIMATION OF THE TRUE POSTERIOR

Based on Proposition 1, the log-posterior probability $\log p_{\theta, w}(\mathbf{z}_I | \mathcal{G}, \mathbf{q}, \mathbf{a})$ could be approximated by $(\sum_{rule \in \mathbf{z}_I} H(rule) + \gamma(\mathbf{z}_I) + \text{const})$, with const being a term that does not depend on \mathbf{z}_I . This implies that we could construct a distribution $q(\mathbf{z}_I) \propto \exp(\sum_{rule \in \mathbf{z}_I} H(rule) + \gamma(\mathbf{z}_I))$ to approximate the true posterior, and draw samples from q as approximation to the real samples from the posterior.

It turns out that the distribution $q(\mathbf{z}_I)$ is a multinomial distribution. To see that, we rewrite $q(\mathbf{z}_I)$ as:

$$\begin{aligned}
q(\mathbf{z}_I) &= \frac{1}{Z} \exp \left(\sum_{rule \in \mathbf{z}_I} H(rule) + \gamma(\mathbf{z}_I) \right) \\
&= \frac{1}{Z} \exp(\gamma(\mathbf{z}_I)) \prod_{rule \in \mathbf{z}_I} \exp(H(rule)) \\
&= \frac{1}{Z} \frac{K!}{\prod_{rule \in \hat{\mathbf{z}}} n_{rule}!} \prod_{rule \in \hat{\mathbf{z}}} \exp(H(rule))^{n_{rule}} \\
&= \frac{1}{Z'} \frac{K!}{\prod_{rule \in \hat{\mathbf{z}}} n_{rule}!} \prod_{rule \in \hat{\mathbf{z}}} q_r(rule)^{n_{rule}} \\
&= \frac{1}{Z'} \text{Mu}(\mathbf{z}_I | K, q_r),
\end{aligned}$$

where n_{rule} is the number of times a *rule* appears in the set \mathbf{z}_I , q_r is a distribution over all the generated logic rules $\hat{\mathbf{z}}$ with $q_r(rule) = \exp(H(rule)) / \sum_{rule' \in \hat{\mathbf{z}}} \exp(H(rule'))$, Z and Z' are normalization terms. By summing over \mathbf{z}_I on both sides of the above equation, we obtain $Z' = 1$, and hence:

$$q(\mathbf{z}_I) = \text{Mu}(\mathbf{z}_I | K, q_r).$$

To sample from such a multinomial distribution, we could simply sample K rules independently from the distribution q_r , and form a sample \hat{z}_I with these K rules.

In practice, we observe that the hard-assignment EM algorithm (Koller & Friedman, 2009) works better than the standard EM algorithm despite the reduced theoretical guarantees. In the hard-assignment EM algorithm, we need to draw a sample \hat{z}_I with the maximum posterior probability. Based on the above approximation $q(z_I)$ of the true posterior distribution $p_{\theta,w}(z_I|\mathcal{G}, \mathbf{q}, \mathbf{a})$, we could simply construct such a sample \hat{z}_I with K rules which have the maximum probability under the distribution q_r . By definition, we have $q_r(\text{rule}) \propto \exp(H(\text{rule}))$, and hence drawing K rules with maximum probability under q_r is equivalent to choosing K rules with the maximum H values.

B MORE ANALYSIS OF THE EM ALGORITHM

In RNNLogic, we use an EM algorithm to optimize the rule generator. In this section, we show why this EM algorithm is able to maximize the objective function of the rule generator.

Recall that for a fixed reasoning predictor p_w , we aim to update p_θ to maximize the log-likelihood function $\log p_{w,\theta}(\mathbf{a}|\mathcal{G}, \mathbf{q})$ for each data instance $(\mathcal{G}, \mathbf{q}, \mathbf{a})$. Directly optimizing $\log p_{w,\theta}(\mathbf{a}|\mathcal{G}, \mathbf{q})$ is difficult due to the latent logic rules, and therefore we consider the following evidence lower bound of the log-likelihood function:

$$\log p_{w,\theta}(\mathbf{a}|\mathcal{G}, \mathbf{q}) \geq \mathbb{E}_{q(z_I)}[\log p_w(\mathbf{a}|\mathcal{G}, \mathbf{q}, z_I) + \log p_\theta(z_I|\mathbf{q}) - \log q(z_I)] = \mathcal{L}_{\text{ELBO}}(q, p_\theta), \quad (12)$$

where $q(z_I)$ is a variational distribution, and the equation holds when $q(z_I) = p_{\theta,w}(z_I|\mathcal{G}, \mathbf{q}, \mathbf{a})$.

With this lower bound, we can optimize the log-likelihood function $\log p_{w,\theta}(\mathbf{a}|\mathcal{G}, \mathbf{q})$ with an E-step and an M-step. In the E-step, we optimize $q(z_I)$ to maximize $\mathcal{L}_{\text{ELBO}}(q, p_\theta)$, which is equivalent to minimizing $\text{KL}(q(z_I)||p_{\theta,w}(z_I|\mathcal{G}, \mathbf{q}, \mathbf{a}))$. By doing so, we are able to tighten the lower bound. Then in the M-step, we further optimize θ to maximize $\mathcal{L}_{\text{ELBO}}(q, p_\theta)$. Next, we introduce the details.

E-step. In the E-step, our goal is to update q to minimize $\text{KL}(q(z_I)||p_{\theta,w}(z_I|\mathcal{G}, \mathbf{q}, \mathbf{a}))$. However, there are a huge number of possible logic rules, and hence optimizing $q(z_I)$ on every possible rule set z_I is intractable. To solve the problem, recall that we generate a set of logic rules \hat{z} when optimizing the reasoning predictor, and here we add a constraint to q based on \hat{z} . Specifically, we constrain the sample space of $q(z_I)$ to be all subsets of \hat{z} with size being K , i.e., $z_I \subset \hat{z}$ and $|z_I| = K$. In other words, we require $\sum_{z_I \subset \hat{z}, |z_I|=K} q(z_I) = 1$. With such a constraint, we can further use proposition 1 to construct the variational distribution q to approximate $p_{\theta,w}(z_I|\mathcal{G}, \mathbf{q}, \mathbf{a})$, as what is described in the model section.

M-step. In the M-step, our goal is to update p_θ to maximize the lower bound $\mathcal{L}_{\text{ELBO}}(q, p_\theta)$. To do that, we notice that there is an expectation operation with respect to $q(z_I)$ in $\mathcal{L}_{\text{ELBO}}(q, p_\theta)$. By drawing a sample from $q(z_I)$, $\mathcal{L}_{\text{ELBO}}(q, p_\theta)$ can be estimated as follows:

$$\begin{aligned} \mathcal{L}_{\text{ELBO}}(q, p_\theta) &= \mathbb{E}_{q(z_I)}[\log p_w(\mathbf{a}|\mathcal{G}, \mathbf{q}, z_I) + \log p_\theta(z_I|\mathbf{q}) - \log q(z_I)] \\ &\simeq \log p_w(\mathbf{a}|\mathcal{G}, \mathbf{q}, \hat{z}_I) + \log p_\theta(\hat{z}_I|\mathbf{q}) - \log q(\hat{z}_I), \end{aligned} \quad (13)$$

where $\hat{z}_I \sim q(z_I)$ is a sample drawn from the variational distribution. By ignoring the terms which are irrelevant to p_θ , we obtain the following objective function for θ :

$$\log p_\theta(\hat{z}_I|\mathbf{q}), \quad (14)$$

which is the same as the objective function described in the model section.

As a result, by performing the E-step and the M-step described in the model section, we are able to update p_θ to increase the lower bound $\mathcal{L}_{\text{ELBO}}(q, p_\theta)$, and thereby push up the log-likelihood function $\log p_{w,\theta}(\mathbf{a}|\mathcal{G}, \mathbf{q})$. Therefore, we see that the EM algorithm can indeed maximize $\log p_{w,\theta}(\mathbf{a}|\mathcal{G}, \mathbf{q})$ with respect to p_θ .

C DETAILS ABOUT PARAMETERIZATION AND IMPLEMENTATION

Section 3.2 of the paper introduces the high-level idea of the reasoning predictor with logic rules and the rule generator. Due to the limited space, some details of the models are not covered. In this section, we explain the details of the reasoning predictor and the rule generator.

C.1 REASONING PREDICTOR WITH LOGIC RULES

We start with the reasoning predictor with logic rules. Recall that for each query, our reasoning predictor leverages a set of logic rules \mathcal{z} to give each candidate answer a score, which is further used to predict the correct answer from all candidates.

Specifically, let \mathcal{A} denote the set of all the candidate answers discovered by logic rules in set \mathcal{z} . For each candidate answer $e \in \mathcal{A}$, we define the following function score_w to compute a score:

$$\text{score}_w(e) = \sum_{rule \in \mathcal{z}} \text{score}_w(e|rule) = \sum_{rule \in \mathcal{z}} \sum_{path \in \mathcal{P}(h, rule, e)} \psi_w(rule) \cdot \phi_w(path), \quad (15)$$

where $\mathcal{P}(h, rule, e)$ is the set of grounding paths which start at h and end at e following a *rule* (e.g., *Alice* $\xrightarrow{\text{friend}}$ *Bob* $\xrightarrow{\text{hobby}}$ *Sing*). $\psi_w(rule)$ and $\phi_w(path)$ are scalar weights of each *rule* and *path*.

For the scalar weight $\psi_w(rule)$ of a *rule*, we initialize $\psi_w(rule)$ as follows:

$$\psi_w(rule) = \mathbb{E}_{(\mathcal{G}, \mathbf{q}, \mathbf{a}) \sim p_{\text{data}}} \left[|\mathcal{P}(h, rule, t)| - \frac{1}{|\mathcal{A}|} \sum_{e \in \mathcal{A}} |\mathcal{P}(h, rule, e)| \right], \quad (16)$$

where $|\mathcal{P}(h, rule, t)|$ is the number of relational paths starting from head entity h , following the relations in *rule* and ending at tail entity t . The form is very similar to the definition of H values for logic rules, and the value can effectively measure the contribution of a rule to the correct answers. We also try randomly initializing rule weights or initializing them as 0, which yield similar results.

For the scalar score $\phi_w(path)$ of a *path*, we either fix it to 1, or compute it by introducing entity and relation embeddings. In the second case, we introduce an embedding for each entity e is denoted as \mathbf{x}_e , and the embedding of a relation \mathcal{r} is denoted as $\mathbf{x}_{\mathcal{r}}$. For a grounding path $path = e_0 \xrightarrow{\mathcal{r}_1} e_1 \xrightarrow{\mathcal{r}_2} e_2 \cdots e_{l-1} \xrightarrow{\mathcal{r}_l} e_l$, we follow the idea in RotatE (Sun et al., 2019) and compute $\phi_w(path)$ in the following way:

$$\phi_w(path) = \sigma(\delta - d(\mathbf{x}_{e_0} \circ \mathbf{x}_{\mathcal{r}_1} \circ \mathbf{x}_{\mathcal{r}_2} \circ \cdots \circ \mathbf{x}_{\mathcal{r}_l}, \mathbf{x}_{e_l})), \quad (17)$$

where $\sigma(x) = \frac{1}{1+e^{-x}}$ is the sigmoid function, d is a distance function between two complex vectors, δ is a hyperparameter, and \circ is the Hadmard product in complex spaces, which could be viewed as a rotation operator. Intuitively, for the embedding \mathbf{x}_{e_0} of entity e_0 , we rotate \mathbf{x}_{e_0} by using the rotation operators defined by $\{\mathcal{r}_k\}_{k=1}^l$, yielding $(\mathbf{x}_{e_0} \circ \mathbf{x}_{\mathcal{r}_1} \circ \mathbf{x}_{\mathcal{r}_2} \circ \cdots \circ \mathbf{x}_{\mathcal{r}_l})$. Then we compute the distance between the new embedding and the embedding \mathbf{x}_{e_l} of entity e_l , and further convert the distance to a value between 0 and 1 by using the sigmoid function and a hyperparameter δ .

C.2 RULE GENERATOR

This paper focuses on compositional rules, which have the abbreviation form $\mathcal{r} \leftarrow \mathcal{r}_1 \wedge \cdots \wedge \mathcal{r}_l$ and thus could be viewed a sequence of relations $[\mathcal{r}, \mathcal{r}_1, \mathcal{r}_2 \cdots \mathcal{r}_l, \mathcal{r}_{\text{END}}]$, where \mathcal{r} is the query relation or the head of the rule, $\{\mathcal{r}_i\}_{i=1}^l$ are the body of the rule, and \mathcal{r}_{END} is a special relation indicating the end of the relation sequence. We introduce a rule generator RNN_θ parameterized with an LSTM (Hochreiter & Schmidhuber, 1997) to model such sequences. Given the current relation sequence $[\mathcal{r}, \mathcal{r}_1, \mathcal{r}_2 \cdots \mathcal{r}_i]$, RNN_θ aims to generate the next relation \mathcal{r}_{i+1} and meanwhile output the probability of \mathcal{r}_{i+1} . The detailed computational process towards the goal is summarized as follows:

- Initialize the hidden state of the RNN_θ as follows:

$$\mathbf{h}_0 = f(\mathbf{v}_{\mathcal{r}}),$$

where $\mathbf{v}_{\mathcal{r}}$ is a parameter vector associated with the query relation or the head relation \mathcal{r} , f is a linear transformation.

- Sequentially compute the hidden state at different positions by using the LSTM gate:

$$\mathbf{h}_t = \text{LSTM}(\mathbf{h}_{t-1}, g([\mathbf{v}_{\mathcal{r}}, \mathbf{v}_{\mathcal{r}_t}]),$$

where $\mathbf{v}_{\mathcal{r}_t}$ is a parameter vector associated with the relation \mathcal{r}_t , $[\mathbf{v}_{\mathcal{r}}, \mathbf{v}_{\mathcal{r}_t}]$ is the concatenation of $\mathbf{v}_{\mathcal{r}}$ and $\mathbf{v}_{\mathcal{r}_t}$, g is a linear transformation.

- Generate r_{t+1} and its probability based on \mathbf{h}_{t+1} and the following vector:

$$\text{softmax}(o(\mathbf{h}_{t+1})).$$

Suppose the set of relations is denoted as \mathcal{R} . We first transform \mathbf{h}_{t+1} to a $|\mathcal{R}|$ -dimensional vector by using a linear transformation o , and then apply softmax function to the $|\mathcal{R}|$ -dimensional vector to get the probability of each relation. Finally, we generate r_{t+1} according to the probability vector.

D EXPERIMENTAL SETUP

D.1 DATASET STATISTICS

The statistics of datasets are summarised in Table 6.

Table 6: Statistics of datasets.

Dataset	#Entities	#Relations	#Train	#Validation	#Test
FB15K-237	14,541	237	272,115	17,535	20,466
WN18RR	40,943	11	86,835	3,034	3,134
Kinship	104	25	3,206	2,137	5,343
UMLS	135	46	1,959	1,306	3,264

D.2 EXPERIMENTAL SETUP OF RNNLOGIC

Next, we explain the detailed experimental setup of RNNLogic. We try different configurations of hyperparameters on the validation set, and the optimal configuration is then used for testing. We report the optimal hyperparameter configuration as below.

Data Preprocessing. For each training triplet (h, r, t) , we add an inverse triplet (t, r^{-1}, h) into the training set, yielding an augmented set of training triplets \mathcal{T} . To build a training instance from p_{data} , we first randomly sample a triplet (h, r, t) from \mathcal{T} , and then form an instance as $(\mathcal{G} = \mathcal{T} \setminus \{(h, r, t)\}, \mathbf{q} = (h, r, ?), \mathbf{a} = t)$. Basically, we use the sampled triplet (h, r, t) to construct the query and answer, and use the rest of triplets in \mathcal{T} to form the background knowledge graph \mathcal{G} . During testing, the background knowledge graph \mathcal{G} is constructed by using all the triplets in \mathcal{T} .

Reasoning Predictor. For the reasoning predictor in *with embedding* cases, the embedding dimension is set to 500 for FB15k-237, 200 for WN18RR, 2000 for Kinship and 1000 for UMLS. We pre-train these embeddings with RotatE (Sun et al., 2019). The hyperparameter δ for computing $\phi_w(\text{path})$ in Equation (17) is set to 9 for FB15k-237, 6 for WN18RR, 0.25 for Kinship and 3 for UMLS. We use the Adam (Kingma & Ba, 2014) optimizer with an initial learning rate being 5×10^{-5} , and we decrease the learning rate in a cosine shape.

Rule Generator. For the rule generator, the maximum length of generated rules is set to 4 for FB15k-237, 5 for WN18RR, and 3 for Kinship and UMLS. These numbers are chosen according to model performance on the validation data. The size of input and hidden states in RNN_θ are set to 512 and 256. The learning rate is set to 1×10^{-3} and monotonically decreased in a cosine shape. Beam search is used to generate rules with high probabilities, so that we focus on exploiting these logic rules which the rule generator is confident about. Besides, we pre-train the rule generator by using sampled relational paths on the background knowledge graph formed with training triplets, which prevents the rule generator from exploring meaningless logic rules in the beginning of training.

EM Optimization. During optimization, we sample 1000 rules from the rule generator for each data instance. In the E-step, for each data instance, we identify 300 rules as high-quality logic rules.

Evaluation. In testing, for each query $\mathbf{q} = (h, r, ?)$, we use the rule generator p_θ to generate 1000 logic rules, and let the reasoning predictor p_w use the generated logic rules to predict the answer \mathbf{a} .

E CASE STUDY

We present more rules learned by RNNLogic on FB15k-237 dataset and UMLS dataset in Table 7 at the last page. In this table, $h \xrightarrow{r} t$ means triplet (h, r, t) and $h \xleftarrow{r} t$ means triplet (h, r^{-1}, t) , or equivalently (t, r, h) .

F CONNECTION WITH REINFORCE

In terms of optimizing the rule generator, our EM algorithm has some connections with the REINFORCE algorithm (Williams, 1992). Formally, for a training instance $(\mathcal{G}, \mathbf{q}, \mathbf{a})$, REINFORCE computes the derivative with respect to the parameters θ of the rule generator as follows:

$$\mathbb{E}_{p_{\theta}(\mathbf{z}|\mathbf{q})}[R \cdot \nabla_{\theta} \log p_{\theta}(\mathbf{z}|\mathbf{q})] \simeq R \cdot \nabla_{\theta} \log p_{\theta}(\hat{\mathbf{z}}|\mathbf{q}), \quad (18)$$

where $\hat{\mathbf{z}}$ is a sample from the rule generator, i.e., $\hat{\mathbf{z}} \sim p_{\theta}(\mathbf{z}|\mathbf{q})$. R is a reward from the reasoning predictor based on the prediction result on the instance. For example, we could treat the probability of the correct answer computed by the reasoning predictor as reward, i.e., $R = p_w(\mathbf{a}|\mathcal{G}, \mathbf{q}, \hat{\mathbf{z}})$.

In contrast, our EM optimization algorithm optimizes the rule generator with the objective function defined in Equation (2), yielding the derivative $\nabla_{\theta} \log p_{\theta}(\hat{\mathbf{z}}_I|\mathbf{q})$. Comparing the derivative to that in REINFORCE (Eq. (18)), we see that EM only maximizes the log-probability for rules selected in the E-step, while REINFORCE maximizes the log-probability for all generated rules weighted by the scalar reward R . Hence the two approaches coincide if R is set to 1 for the selected rules from the approximate posterior and 0 otherwise. In general, finding an effective reward function to provide feedback is nontrivial, and we empirically compare these two optimization algorithms in experiments.

Table 7: Logic rules learned by RNNLogic.

Relation	←	Rule (Explanation)
$X \xrightarrow{\text{Appears.in.TV.Show}} Y$	←	$X \xleftarrow{\text{Has.Actor}} Y$ (Definition. An actor of a show appears in the show, obviously.)
	←	$X \xrightarrow{\text{Creator.of}} U \xleftarrow{\text{Producer.of}} V \xrightarrow{\text{Appears.in.TV.Show}} Y$ (The creator X and the producer V of another show U are likely to appear in the same show Y .)
	←	$X \xleftarrow{\text{Actor.of}} U \xleftarrow{\text{Award.Nominated}} V \xleftarrow{\text{Winner.of}} Y$
	←	$X \xrightarrow{\text{Writer.of}} U \xleftarrow{\text{Creator.of}} V \xrightarrow{\text{Actor.of}} Y$
	←	$X \xrightarrow{\text{Student.of}} U \xleftarrow{\text{Student.of}} V \xrightarrow{\text{Appears.in.TV.Show}} Y$ (Two students X and V in the same school U are likely to appear in the same show Y .)
$X \xrightarrow{\text{ORG.in.State}} Y$	←	$X \xrightarrow{\text{ORG.in.City}} U \xrightarrow{\text{City.in.State}} Y$ (Use the city to indicate the state directly.)
	←	$X \xrightarrow{\text{ORG.in.City}} U \xleftarrow{\text{Lives.in}} V \xrightarrow{\text{Born.in}} W \xrightarrow{\text{Town.in.State}} Y$ (Use the person living in the city to induct the state.)
	←	$X \xleftarrow{\text{Sub-ORG.of}} U \xrightarrow{\text{ORG.in.State}} Y$
	←	$X \xrightarrow{\text{Sub-ORG.of}} U \xleftarrow{\text{Sub-ORG.of}} V \xrightarrow{\text{ORG.in.State}} Y$
	←	$X \xrightarrow{\text{ORG.in.City}} U \xleftarrow{\text{ORG.in.City}} V \xrightarrow{\text{ORG.in.State}} Y$ (Two organizations in the same city are in the same state.)
$X \xrightarrow{\text{Person.Nationality}} Y$	←	$X \xrightarrow{\text{Born.in}} U \xrightarrow{\text{Place.in.Country}} Y$ (Definition.)
	←	$X \xrightarrow{\text{Spouse}} U \xrightarrow{\text{Person.Nationality}} Y$ (By a fact that people are likely to marry a person of same nationality.)
	←	$X \xrightarrow{\text{Student.of}} U \xrightarrow{\text{ORG.Endowment.Currency}} V \xleftarrow{\text{Region.Currency}} W \xrightarrow{\text{Region.in.Country}} Y$ (Use the currency to induct the nationality.)
	←	$X \xrightarrow{\text{Born.in}} U \xleftarrow{\text{Born.in}} V \xrightarrow{\text{Person.Nationality}} Y$
	←	$X \xrightarrow{\text{Politician.of}} U \xleftarrow{\text{Politician.of}} V \xrightarrow{\text{Person.Nationality}} Y$
$X \xrightarrow{\text{Manifestation.of}} Y$	←	$X \xleftarrow{\text{Treats}} U \xrightarrow{\text{Prevents}} V \xleftarrow{\text{Precedes}} Y$
	←	$X \xleftarrow{\text{Complicates}} U \xleftarrow{\text{Precedes}} Y$
	←	$X \xrightarrow{\text{Location.of}} U \xrightarrow{\text{Is.a}} V \xleftarrow{\text{Precedes}} Y$
	←	$X \xleftarrow{\text{Complicates}} U \xrightarrow{\text{Precedes}} V \xleftarrow{\text{Occurs.in}} Y$
	←	$X \xrightarrow{\text{Location.of}} U \xleftarrow{\text{Occurs.in}} V \xleftarrow{\text{Occurs.in}} Y$
	←	$X \xrightarrow{\text{Precedes}} U \xleftarrow{\text{Occurs.in}} V \xleftarrow{\text{Degree.of}} Y$
$X \xleftarrow{\text{Affects}} Y$	←	$X \xrightarrow{\text{Result.of}} U \xrightarrow{\text{Occurs.in}} V \xrightarrow{\text{Precedes}} Y$
	←	$X \xleftarrow{\text{Precedes}} U \xrightarrow{\text{Produces}} V \xleftarrow{\text{Occurs.in}} Y$
	←	$X \xleftarrow{\text{Prevents}} U \xrightarrow{\text{Disrupts}} V \xrightarrow{\text{Co-occurs.with}} Y$
	←	$X \xleftarrow{\text{Result.of}} U \xrightarrow{\text{Complicates}} V \xrightarrow{\text{Precedes}} Y$
	←	$X \xleftarrow{\text{Assesses.Effect.of}} U \xrightarrow{\text{Method.of}} V \xrightarrow{\text{Complicates}} Y$
	←	$X \xrightarrow{\text{Process.of}} U \xrightarrow{\text{Interacts.with}} V \xrightarrow{\text{Causes}} Y$
	←	$X \xleftarrow{\text{Assesses.Effect.of}} U \xleftarrow{\text{Result.of}} V \xrightarrow{\text{Precedes}} Y$