

Learning Collaborative Agents with Rule Guidance for Knowledge Graph Reasoning

Deren Lei^{1*}, Gangrong Jiang^{1*}, Xiaotao Gu², Kexuan Sun¹, Yuning Mao², Xiang Ren¹

¹University of Southern California

²University of Illinois at Urbana-Champaign

{derenlei, gjiang, kexuansu, xiangren}@usc.edu, {xiaotao2, yuningm2}@illinois.edu

Abstract

Walk-based models have shown their advantages in knowledge graph (KG) reasoning by achieving decent performance while providing interpretable decisions. However, the sparse reward signals offered by the KG during traversal are often insufficient to guide a sophisticated walk-based reinforcement learning (RL) model. An alternate approach is to use traditional symbolic methods (e.g., rule induction), which achieve good performance but can be hard to generalize due to the limitation of symbolic representation. In this paper, we propose *RuleGuider*, which leverages high-quality rules generated by symbolic-based methods to provide reward supervision for walk-based agents. Experiments on benchmark datasets show that *RuleGuider* improves the performance of walk-based models without losing interpretability. ¹

1 Introduction

While knowledge graphs (KGs) are widely adopted in natural language processing applications, a major bottleneck hindering its usage is the sparsity of facts (Min et al., 2013), leading to extensive studies on KG completion (or reasoning) (Trouillon et al., 2016; Dettmers et al., 2018; Das et al., 2017; Xiong et al., 2017; Lin et al., 2018; Meilicke et al., 2019). Many traditional approaches on the KG reasoning task are based on logic rules (Landwehr et al., 2007, 2010; Galárraga et al., 2013, 2015). These methods are referred to as *symbolic-based methods*. Although they showed good performance (Meilicke et al., 2019, 2020), they are inherently limited by their representations and generalizability of the associated relations of the given rules.

To ameliorate such limitations, *embedding-based methods* (Bordes et al., 2013; Socher et al.,

2013; Wang et al., 2014; Yang et al., 2014; Trouillon et al., 2016; Dettmers et al., 2018, 2017; Sun et al., 2019; Zhang et al., 2019) were proposed. They learn distributed representations for entities and relations and make predictions using the representations. Despite their superior performance, they fail to make human-friendly interpretations.

To improve the interpretability, many recent efforts formulate the task as a multi-hop reasoning problem using reinforcement learning (RL) techniques (Xiong et al., 2017; Das et al., 2017; Shen et al., 2018; Chen et al., 2018; Lin et al., 2018), referred to as *walk-based methods*. A major issue of these methods is the reward function. A “hit or not” reward is too sparse while a shaped reward using an embedding-based distance measurement Lin et al. (2018) may not always result in desirable paths.

In this paper, we propose *RuleGuider* to tackle the aforementioned reward issue in walk-based methods with the help of symbolic rules. We aim to improve the performance of walk-based methods without losing their interpretability. The *RuleGuider* is composed of a symbolic-based model fetching logic rules and a walk-based agent searching reasoning paths with the guidance of the rules. We also introduce a way to separate the walk-based agent to allow for further efficiency. We experimentally show the efficiency of our model without losing the interpretability.

2 Problem and Preliminaries

In this section, we review the KG reasoning task. We also describe the symbolic-based and walk-based methods used in *RuleGuider*.

Problem Formulation. A KG consisting of fact triples is represented as $\mathcal{G} = \{(e_i, r, e_j)\} \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{E}$, where \mathcal{E} and \mathcal{R} are the set of entities and relations, respectively. Given a query $(e_s, r_q, ?)$ where e_s is a subject entity and r_q is a query re-

*Equal contributions.

¹<https://github.com/derenlei/KG-RuleGuider>




Figure 1: Rule quality difference between datasets. There are exists high quality rules on WN18RR.

lation, the task of KG reasoning is to find a set of object entities E_o such that (e_s, r_q, e_o) , where $e_o \in E_o$, is a fact triple missing in \mathcal{G} . We denote the queries $(e_s, r_q, ?)$ as *tail queries*. We note that we can also perform *head queries* $(?, r_q, e_o)$. To be consistent with most existing works, we only consider tail queries in this paper.

Symbolic-based Methods. Some previous methods mine Horn rules from the KG and predict missing facts by grounding these rules. A recent method AnyBURL (Meilicke et al., 2019) showed comparable performance to the state-of-the-art embedding-based methods. It first mines rules by sampling paths from the \mathcal{G} , and then make predictions by matching queries to the rules. Rules are in the format: $r(X, Y) \leftarrow b_1(X, A_2) \wedge \dots \wedge b_n(A_n, Y)$, where upper-case letters represent variables. A *rule head* is denoted by $r(\dots)$ and a *rule body* is denoted by the conjunction of atoms $b_1(\dots), \dots, b_n(\dots)$. We note that $r(c_i, c_j)$ is equivalent to the fact triple (c_i, r, c_j) .

However, these methods have limitations. For example, rules mined from different KGs may have different qualities, which makes the reasoner hard to select rules. Figure 1 shows such difference. Rules are sorted based on accuracy of predicting the target entities. The top rules from WN18RR are much more valuable than those from FB15K-237.

Walk-based Methods. Given a query $(e_s, r_q, ?)$, walk-based methods train an RL agent to find a path from e_s to the desired object entity e_o that implies the query relation r_q . At step t , the current state is represented by a tuple $s_t = (e_t, (e_s, r_q))$, where e_t is the current entity. The agent then samples the next relation-entity pair to visit from possible actions $A_t = \{(r', e') | (e_t, r', e') \in \mathcal{G}\}$. The agent receives a reward when it reaches e_o .

3 Proposed Method: RuleGuider

RuleGuider consists of a symbolic-based method (see Section 2), referred to as *rule miner*, and a walk-based method, referred to as *agent*. The rule

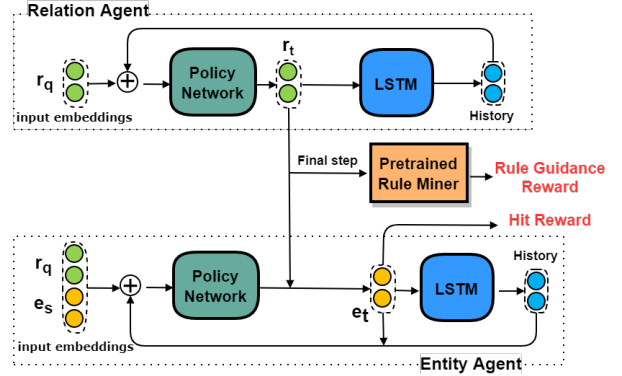


Figure 2: **The architecture of two agents.** The relation and entity agent interact with each other to generate a path. At each step, the entity agent first selects an entity from valid entities. The relation agent then samples a relation based on the selected entity. At the final step, they receive a hit reward based on the last selected entity and a rule guidance reward from the pre-mined rule set based on the selected path.

miner first mines logic rules and the agent traverses over the KG to learn the probability distribution of reasoning paths with the guidance (via the reward) of the rules. As the agent walks through relations and entities alternatively, we propose to separate the agent into two sub-agents: a relation and entity agents. After the separation, the search space is significantly pruned. Figure 2 shows the structure of these two agents in detail.

3.1 Model Architecture

Relation Agent. At step t ($t = 1, \dots, T$, T is the number of hops), the relation agent selects a single relation r_t which is incident to the current entity e_{t-1} , where $e_0 = e_s$. Given a query $(e_s, r_q, ?)$ and a set of rules \mathbf{R} , this process can be formulated as $r_t = P^R(r_q, e_{t-1}, \mathbf{R}, \mathbf{h}_t^R)$ where \mathbf{h}_t^R is the relation history. The agent first filter out rules whose heads are not same as r_q , and then it selects r_t from the t^{th} atoms of the remaining rule bodies, i.e. $b_t(\dots)$ in the rule pattern.

Since the rule miner provides confidence scores of rules, we first use RL techniques to pre-train this agent using the scores. During training, the agent applies the pre-trained strategy (distribution) and keeps tuning the distribution by utilizing semantic information provided by embeddings. In another words, the relation agent leverages both confidence scores of pre-mined rules as well as embedding shaped hit rewards.

Entity Agent. At step t , the agent generates the distribution of all candidate entities based on e_s ,

r_q , and the entity history \mathbf{h}_t^E . Given the current relation r_t , this process can formally be represented as $e_t = P^E(e_s, r_q, r_t, \mathbf{h}_t^E)$. The agent selects an entity from all entities that incident on r_t . In this way, the entity and relation agent can reason independently.

In experiments, we have also tried to let the entity agent generate distribution based on relation agent pruned entity space. In this way, the entity agent takes in the selected relation and can leverage the information from the relation agent. However, the entity space may be extremely small and hard to learn. It makes the entity agent less effective, especially on large and dense KG.

Policy Network. The relation agent’s search policy is parameterized by the embedding of r_q and \mathbf{h}_t^R . The relation history is encoded using an LSTM(Hochreiter and Schmidhuber, 1997): $\mathbf{h}_t^R = \text{LSTM}(\mathbf{h}_{t-1}^R, r_{t-1})$, where $r_{t-1} \in \mathbb{R}^d$ is the embedding of the last relation. We initialize $\mathbf{h}_0^R = \text{LSTM}(\mathbf{0}, r_s)$, where r_s is a special start relation embedding to form an initial relation-entity pair with source entity embedding e_s . Relation space embeddings $\mathbf{R}_t \in \mathbb{R}^{|R_t| \times d}$ consist embeddings of all the relations in relation space R_t at step t . Finally, relation agent outputs a probability distribution \mathbf{d}_t^R and samples a relation from it. $\mathbf{d}_t^R = \sigma(\mathbf{R}_t \times \mathbf{W}_1 \text{ReLU}(\mathbf{W}_2[\mathbf{h}_t^R; r_q]))$ where σ is the softmax operator, \mathbf{W}_1 and \mathbf{W}_2 is trainable parameters. We design relation agent’s history-dependent policy as $\pi^R = (\mathbf{d}_1^R, \mathbf{d}_2^R, \dots, \mathbf{d}_T^R)$.

Similarly, entity agent’s history-dependent policy is $\pi^E = (\mathbf{d}_1^E, \mathbf{d}_2^E, \dots, \mathbf{d}_T^E)$. Entity agent can acquire its embedding of last step e_{t-1} , entity space embeddings \mathbf{E}_t , its history $\mathbf{h}_t^E = \text{LSTM}(\mathbf{h}_{t-1}^E, e_{t-1})$, and the probability distribution of entities \mathbf{d}_t^E as follows. $\mathbf{d}_t^E = \sigma(\mathbf{E}_t \times \mathbf{W}_3 \text{ReLU}(\mathbf{W}_4[\mathbf{h}_t^E; r_q; e_s; e_t]))$ where \mathbf{W}_3 and \mathbf{W}_4 is trainable parameters. Note that entity agent uses a different LSTM to encode the entity history.

3.2 Model Learning

We train the model by letting the two aforementioned agents to start from specific entities and traverse through the KG in a fixed number of hops. The agents receive rewards at their final step.

Reward Design. Given a query, the relation agent prefers paths which direct the way to the correct object entity. Thus, given a relation path, we give reward according to its confidence retrieved from the rule miner, referred to as *rule guidance reward*

R_r . We also add a Laplace smoothing $p_c = 5$ to the confidence score for the final R_r .

In addition to R_r , the agent will also receive a *hit reward* R_h , which is 1 if the predicted triple $\epsilon = (e_s, r_q, e_T) \in \mathcal{G}$. Otherwise, we use the embedding of ϵ to measure reward as in Lin et al. (2018). $R_h = \mathbb{I}(\epsilon \in \mathcal{G}) + (1 - \mathbb{I}(\epsilon \in \mathcal{G}))f(\epsilon)$ where $\mathbb{I}(\cdot)$ is an indicator function, $f(\epsilon)$ is a composition function for reward shaping using embeddings.

Training Procedure. We train the model in four stages. 1) Train relation and entity embeddings using an embedding-based method. 2) Apply a rule miner to retrieve rules and their associated confidence scores. 3) Pre-train the relation agent by freezing the entity agent and asking the relation agent to sample a path. We only use the rule miner to evaluate the path and compute R_r based on the pre-mined confidence score. 4) Jointly train the relation and entity agent to leverage the embeddings to compute R_h . The final reward R involves R_r and R_h with a constant factor λ : $R = \lambda R_r + (1 - \lambda)R_h$. The policy networks of two agents are trained using the REINFORCE (Williams, 1992) algorithm to maximize R .

4 Experiments

In this section, we compare RuleGuider with other approaches on three datasets. We describe the experiment setting, results, and analysis.

4.1 Experimental Setup

Datasets. We evaluate different methods on three benchmark datasets. (1) FB15k-237 (Toutanova et al., 2015), (2) WN18RR (Dettmers et al., 2018), and (3) NELL-995 (Xiong et al., 2017).

Hyperparameters. We set all embedding size to 200. Each history encoder is a three-layer LSTM with a hidden state dimension 200. We use AnyBURL (Meilicke et al., 2019) as the rule miner and set the confidence score threshold to be 0.15. Other hyperparameters are shown in appendix.

4.2 Results

Table 1 shows the evaluation results. RuleGuider achieves the state-of-the-art results over walk-based methods on WN18RR and NELL-995, and also competitive results on FB15k-237. One possible reason is: compared to the other two datasets, the relation space in FB15k-237 is much larger and the rules is relatively sparse in the large relational

Method / Dataset	WN18RR				NELL-995				FB15k-237				
	H@1	H@5	H@10	MRR	H@1	H@5	H@10	MRR	H@1	H@5	H@10	MRR	
Interpretable	MINERVA (Das et al., 2017)	41.3	-	51.3	44.8	66.3	-	83.1	72.5	21.7	-	45.6	29.3
	MultiHop (ConvE) (Lin et al., 2018)	41.4	48.1	51.7	44.8	65.6	-	84.4	72.7	32.7	-	56.4	40.7
	MultiHop (ComplEx) (Lin et al., 2018)	42.5	49.4	52.6	46.1	64.4	79.1	81.6	71.2	32.9	-	54.4	39.3
	AnyBURL (C rules) (Meilicke et al., 2019)	42.9	51.6	53.7	-	44.0	56.0	57.0	-	26.9	43.1	52.0	-
	RuleGuider (ConvE)	42.2	49.9	53.6	46.0	66.0	82.0	85.1	73.1	31.6	49.6	57.4	40.8
	RuleGuider (ComplEx)	44.3	52.4	55.5	48.0	66.4	82.7	85.9	73.6	31.3	49.2	56.4	39.5
Embedding	DistMult (Yang et al., 2014)	35.7	-	38.4	36.7	55.2	-	78.3	64.1	32.4	-	60.0	41.7
	ComplEx (Trouillon et al., 2016)	41.5	45.6	46.9	43.4	63.9	81.7	84.8	72.1	33.7	54.0	<u>62.4</u>	43.2
	ConvE (Dettmers et al., 2018)	40.1	49.8	53.7	44.6	<u>66.7</u>	<u>85.3</u>	<u>87.2</u>	<u>75.1</u>	<u>34.1</u>	<u>54.7</u>	62.2	<u>43.5</u>
	RotateE (Sun et al., 2019)	<u>42.2</u>	<u>51.3</u>	<u>54.1</u>	<u>46.4</u>	-	-	-	-	32.2	53.2	61.6	42.2

Table 1: **Performance comparison with walk-based approaches.** Best scores among the interpretable methods and embedding-based methods are bold and underlined, respectively. In addition, we present the reported scores for state-of-the-art embedding-based methods as reference. We underscore the best performing ones in this category.

Phase	WN18RR	NELL-995	FB15K-237
Pre-training	69.2%	44.9%	46.1%
Training	40.7%	24.5%	41.5%

Table 2: Percentage of rules used by RuleGuider (ComplEx) to predict e_o (beam 0) during inference on the development set at the end of pre-training and training phase.

path space, which makes it harder for the relation agent to select a desired rule.

We find symbolic-based models perform strongly on WN18RR and we set λ higher to receive stronger rule guidance. We also observe that embedding-based methods have consistently good performance on all datasets compared to walk-based methods despite their simplicity. One possible reason is that embedding-based methods implicitly encode the connectivity of the whole graph into the embedding space (Lin et al., 2018). Embedding-based methods are free from the strict traversing in the graph and sometimes benefit from this property due to incompleteness of the graph. By leveraging rules, we also incorporate some global information as guidance to make up for the potential searching space loss during the discrete inference process.

Table 2 shows the percentage of rules used on the development set using ComplEx embedding in the pre-training and training phase. It shows that our model abandons a few rules to further improve hit performance during the training phase.

4.3 Ablation Study

We run different variants of RuleGuider on the development set of WN18RR. We use ComplEx hit reward shaping for consistency. Table 3 shows

Model	Freeze	No	Single	Ours
H@1	41.9	41.8	42.4	42.9
MRR	45.5	45.7	46.4	46.5

Table 3: **Freeze**, **No** and **Single** represent models with freezing pre-trained relation agent, without pre-training and without separating the agent.

	MultiHop	Tie	RuleGuider
Vote	36.92%	0%	63.08%

Table 4: Human evaluation vote between MultiHop (Lin et al., 2018) and ruleGuider for correctly predicted path on FB15K-237 development set. Both model use ComplEx reward shaping.

the results. Freezing pre-trained agent performing worse indicates that hit reward is necessary. Removing pre-training performing worse shows that the walk-based agents benefit from logic rules. The single agent variant performing worse shows the effectiveness of pruning action space.

4.4 Human Evaluation

Besides the evaluation metrics, we further analyze whether the reasoning path that leads to correctly predicted entity is reasonable. We perform human evaluation on the Amazon Mechanical Turk. We randomly sample a evaluation set with 300 triples from development set using uniform distribution on FB15k-237. During evaluation, given the ground truth triple, three judges are asked to choose which path is a better explanation/decomposition of it between: 1. path generated by our method. 2. paths generated by MultiHop’s method. 3. Draw or none of them are reasonable. Note that there are 2.6% of the predicted paths are the same and

they are excluded from the evaluation set. For each triple, we count majority vote as the evaluation result. As it's possible that the three judges each choose a different option which leads to one vote on each choice. In this case, we do not count it in the final evaluation result (table 4). Unexpectedly, no triple get more than one vote on Tie. RuleGuider achieves a better performance and the reasoning path makes more sense to human judges comparing to Multihop with ComplEx reward shaping.

5 Conclusions

In this paper, we proposed a collaborative framework utilizing both symbolic-based and walk-based models. We separate the walk-based agent into an entity and relation agent to effectively leverage the symbolic rules and significantly reduce the action space. Experimentally, our approach improved the performance of the state-of-the-art walk-based models on two benchmark KGs.

In future work, we would like to study how to introduce acyclic rules to the walk-based systems.

References

- Antoine Bordes, Nicolas Usunier, Alberto García-Durán, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *NIPS*.
- Wenhu Chen, Wenhan Xiong, Xifeng Yan, and William Yang Wang. 2018. Variational knowledge graph reasoning. In *Proceedings of NAACL-HLT*, pages 1823–1832.
- Rajarshi Das, Shehzaad Dhuliawala, Manzil Zaheer, Luke Vilnis, Ishan Durugkar, Akshay Krishnamurthy, Alex Smola, and Andrew McCallum. 2017. Go for a walk and arrive at the answer: Reasoning over paths in knowledge bases using reinforcement learning. *arXiv preprint arXiv:1711.05851*.
- Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. 2017. Convolutional 2d knowledge graph embeddings. In *AAAI*.
- Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. 2018. Convolutional 2d knowledge graph embeddings. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Luis Galárraga, Christina Teflioudi, Katja Hose, and Fabian M Suchanek. 2015. Fast rule mining in ontological knowledge bases with amie++. *The VLDB Journal—The International Journal on Very Large Data Bases*, 24(6):707–730.
- Luis Antonio Galárraga, Christina Teflioudi, Katja Hose, and Fabian Suchanek. 2013. Amie: association rule mining under incomplete evidence in ontological knowledge bases. In *Proceedings of the 22nd international conference on World Wide Web*, pages 413–422. ACM.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Niels Landwehr, Kristian Kersting, and Luc De Raedt. 2007. Integrating naive bayes and foil. *Journal of Machine Learning Research*, 8(Mar):481–507.
- Niels Landwehr, Andrea Passerini, Luc De Raedt, and Paolo Frasconi. 2010. Fast learning of relational kernels. *Machine learning*, 78(3):305–342.
- Xi Victoria Lin, Richard Socher, and Caiming Xiong. 2018. Multi-hop knowledge graph reasoning with reward shaping. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3243–3253, Brussels, Belgium. Association for Computational Linguistics.
- Christian Meilicke, Melisachew Wudage Chekol, Manuel Fink, and Heiner Stuckenschmidt. 2020. Reinforced anytime bottom up rule learning for knowledge graph completion. *arXiv preprint arXiv:2004.04412*.
- Christian Meilicke, Melisachew Wudage Chekol, Daniel Ruffinelli, and Heiner Stuckenschmidt. 2019. Anytime bottom-up rule learning for knowledge graph completion. In *IJCAI*.
- Bonan Min, Ralph Grishman, Li Wan, Chang Wang, and David Gondek. 2013. Distant supervision for relation extraction with an incomplete knowledge base. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 777–782.
- Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab.
- Yelong Shen, Jianshu Chen, Po-Sen Huang, Yuqing Guo, and Jianfeng Gao. 2018. Reinforcewalk: Learning to walk in graph with monte carlo tree search.
- Richard Socher, Danqi Chen, Christopher D. Manning, and Andrew Y. Ng. 2013. Reasoning with neural tensor networks for knowledge base completion. In *NIPS*.
- Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2019. Rotate: Knowledge graph embedding by relational rotation in complex space. *ArXiv*, abs/1902.10197.

Kristina Toutanova, Danqi Chen, Patrick Pantel, Hoi-fung Poon, Pallavi Choudhury, and Michael Gamon. 2015. Representing text for joint embedding of text and knowledge bases. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1499–1509.

Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. Complex embeddings for simple link prediction. In *ICML*.

Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. Knowledge graph embedding by translating on hyperplanes. In *Twenty-Eighth AAAI conference on artificial intelligence*.

Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.

Wenhan Xiong, Thien Hoang, and William Yang Wang. 2017. DeepPath: A reinforcement learning method for knowledge graph reasoning. *arXiv preprint arXiv:1707.06690*.

Bishan Yang, Wen tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2014. Embedding entities and relations for learning and inference in knowledge bases. *CoRR*, abs/1412.6575.

Wen Zhang, Bibek Paudel, Wei Zhang, Abraham Bernstein, and Huajun Chen. 2019. [Interaction embeddings for prediction and explanation in knowledge graphs](#). *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining - WSDM '19*.

A Appendix

A.1 Datasets

We use the same training, development and testing set splits as [Lin et al. \(2018\)](#). Following [Lin et al. \(2018\)](#), we restrict the output degree of an entity by selecting top η neighbors according to their PageRank score ([Page et al., 1999](#)). We remove unseen entities in test set for NELL-995. We also add reverse links from object entity e_o to subject entity e_s .

The detailed datasets statistics is shown in Table 5 and Table 6.

Dataset	#Ent	#Rel	#Fact
FB15k-237	14,505	237	272,115
WN18RR	40,945	11	86,835
NELL-995	75,492	200	154,213

Table 5: Number of relations, entities and fact triples on three datasets.

Dataset	Degree		Relation Degree		Entity Degree	
	mean	median	mean	median	mean	median
FB15k-237	37.52	22	10.32	10	29.17	18
NELL-995	4.03	1	1.79	1	3.47	1
WN18RR	4.28	3	2.55	2	3.54	2

Table 6: Output degree of each entity on three datasets. Degree is the total edges incident to each entity. Relation degree is the number of relations on each entity’s output edges. Entity degree is the number of entities that each entity connects to.

A.2 Training Details

A.2.1 Hardware and Runtime

We trained our model on one NVIDIA GeForce 1080 Ti GPU. Table 7 shows the runtime detail of our model training.

	WN18RR	NELL-995	FB15k-237
number of epochs	50	1000	30
time/epoch	360s	50s	1800s
trainable parameters	26M	47M	11M

Table 7: Running time and model parameters.

A.2.2 Hyperparameters Search

We use Adam to train our model and use beam search during inference to give the ranked prediction of object entities. We run grid search on a bunch of hyperparameters to select the best configuration. The bounds for searched hyperparameters are shown in Table 8. An exception is rule guidance reward ratio λ , which are manually tuned. Table 10 shows the configurations of our best performing model.

Hyperparameter	Search Bounds
regularization weight β	[0.0, 0.1]
embedding dropout rate	[0.0, 0.3]
hidden layer dropout rate	[0.0, 0.3]
relation dropout rate	[0.0, 0.95]
entity dropout rate	[0.0, 0.95]
bandwidth	{200, 256, 400, 512}
mini-batch size	{64, 128, 256, 512}
learning rate	[0.001, 0.003]
number of hops	{2, 3}

Table 8: Searched hyperparameters using grid search. Following ([Lin et al., 2018](#)), we add an entropy regularization term in the training objective and the term is weighted by regularization weight β . Bandwidth is the entity output degree.

A.2.3 Confidence Score Threshold

We analyze the performance of our model with different confidence score thresholds of the rule miner (AnyBURL). We set the maximum threshold to be 0.15. Table 9 shows the results. The results do

not present any observable pattern. One potential reason is that walk-based reasoning paths and less confident rules may have similar performance on certain queries.

Confidence	> 0.00	> 0.05	> 0.10	> 0.15
@1	42.7	42.7	42.4	42.9
MRR	46.3	46.6	46.0	46.5

Table 9: Different confidence score thresholds.

Hyperparameter	WN18RR	NELL-995	FB15k-237
regularization weight β	0.0	0.05	0.02
embedding dropout rate	0.1	0.1	0.3
hidden layer dropout rate	0.1	0.1	0.1
relation dropout rate	0.1	0.1	0.5
entity dropout rate	0.1	0.1	0.5
bandwidth	500	256	400
mini-batch size	256	128	256
learning rate	0.001	0.001	0.0015
number of hops	3	3	3
rule guidance reward ratio λ	0.65	0.1	0.1

Table 10: Hyperparameter used in our model.

A.2.4 Evaluation Metrics

During inference, the model gives a ranked list of predicted entities as the result. We use Hit@N (H@N) and Mean Reciprocal Rank (MRR) to evaluate the model performance based on the ranked list. Hit@N measures the percentage of test triples for which the correct object entity is ranked in top N in the candidate entities list. MRR measures the average reciprocal of the rank of the object entity.

A.3 Development Set Performance

The result of RuleGuider with complex embedding hit reward shaping on development set is shown in Table 11.

Dataset	H@1	H@5	H@10	MRR
WN18RR	43.7	50.2	53.0	47.0
NELL-995	74.3	92.9	93.4	83.0
FB15K-237	27.4	47.4	55.8	36.7

Table 11: Performance of RuleGuider using ComplEx embedding on development set.