

# Neuro-Symbolic Reinforcement Learning with First-Order Logic

Daiki Kimura Masaki Ono Subhajit Chaudhury Ryosuke Kohita Akifumi Wachi  
Don Joven Agravante Michiaki Tatsubori Asim Munawar Alexander Gray  
IBM Research

{daiki, moonoo, subhajit}@jp.ibm.com, {kohi, akifumi.wachi}@ibm.com  
don.joven.r.agravante@ibm.com, mich@jp.ibm.com, {asim, alexander.gray}@ibm.com

## Abstract

Deep reinforcement learning (RL) methods often require many trials before convergence, and no direct interpretability of trained policies is provided. In order to achieve fast convergence and interpretability for the policy in RL, we propose a novel RL method for text-based games with a recent neuro-symbolic framework called Logical Neural Network, which can learn symbolic and interpretable rules in their differentiable network. The method is first to extract first-order logical facts from text observation and external word meaning network (ConceptNet), then train a policy in the network with directly interpretable logical operators. Our experimental results show RL training with the proposed method converges significantly faster than other state-of-the-art neuro-symbolic methods in a TextWorld benchmark.

## 1 Introduction

Deep reinforcement learning (RL) has been successfully applied to many applications, such as computer games, text-based games, and robot control applications (Mnih et al., 2015; Narasimhan et al., 2015; Kimura, 2018; Yuan et al., 2018; Kimura et al., 2018). However, these methods require many training trials for converging to the optimal action policy, and the trained action policy is not understandable for human operators. This is because, although the training results are sufficient, the policy is stored in a black-box deep neural network. These issues become critical problems when the human operator wants to solve a real-world problem and verify the trained rules. If the trained rules are understandable and modifiable, the human operator can control them and design an action restriction. While using a symbolic (logical) format as representation for stored rules is suitable for achieving interpretability and quick training, it is difficult to train the logical rules with a traditional training approach.

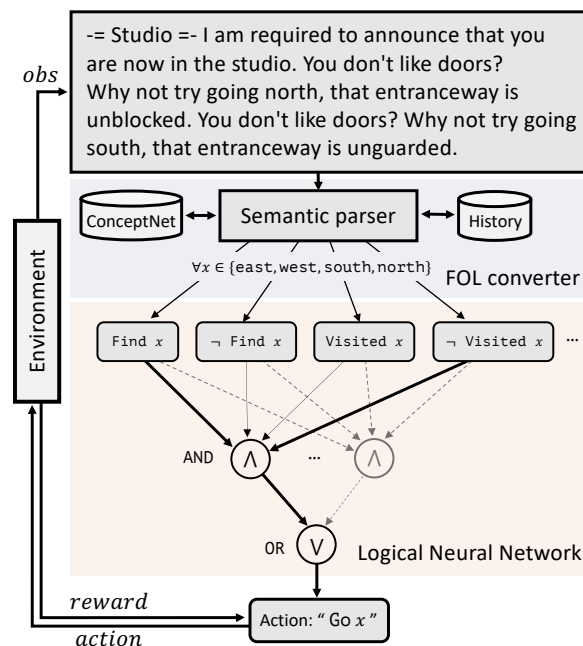


Figure 1: Overview of the proposed method. The agent takes a text observation from the environment, and the first-order logical facts are extracted from an FOL converter that uses a semantic parser, ConceptNet, and history. The weights (shown by line thickness in this figure) of the network are updated by these extracted predicate logics. Solid lines show one trained rule; when the agent finds a direction  $x$  and the direction  $x$  has not been visited, the agent takes a “Go  $x$ ” action. Dashed lines show the initial connections before training.

In order to train logical rules, a recent neuro-symbolic framework called the Logical Neural Network (LNN) (Riegel et al., 2020) has been proposed to simultaneously provide key properties of both the neural network (learning) and the symbolic logic (reasoning). The LNN can train the symbolic rules with logical functions in the neural networks by having an end-to-end differentiable network minimizes a contradiction loss. Every neuron in the LNN has a component for a formula of weighted real-valued logics from a unique logical conjunction, disjunction, or negation nodes, and

then it can calculate the probability and logical contradiction loss during the inference and training. At the same time, the trained LNN can extract obtained logical rules by selecting high weighted connections that represent the important rules for an action policy.

In this paper, we propose an action knowledge acquisition method featuring a neuro-symbolic LNN framework for the RL algorithm. Through experiments, we demonstrate the advantages of the proposed method for real-world problems which is not logically grounded games such as Blocks World. Since natural language observation is easier to convert into logical information than visual or audio, we tackle text-based interaction games for verifying the proposed method.

Figure 1 shows an overview of our method. The observation text is input to a semantic parser to extract the logical values of each propositional logic. In this case, the semantic parser finds there are two exits (north and south). The method then converts first-order logical (predicates) facts from the propositional logics and categories of each word, such as  $\exists x \in \{\text{south, north}\}, \langle \text{find } x \rangle = \text{True}$  and  $\exists x \in \{\text{east, west}\}, \langle \text{find } x \rangle = \text{False}$ . These extracted predicated logics are fed into LNN which has some conjunction gates and one disjunction gate. The LNN trains the weights for these connections by the reward value to obtain the action policy. The contributions of this paper are as follows.

- The paper describes design and implementation of a novel neuro-symbolic RL for a text-based interaction games.
- The paper explains an algorithm to extract first-order logical facts from given textual observation by using the agent history and ConceptNet as an external knowledge.
- We observed our proposed method has advantages for faster convergence and interpretability than state-of-the-art methods and baselines by ablation study on the text-based games.

## 2 Related work

Various prior works have examined RL for text-based games. LSTM-DQN (Narasimhan et al., 2015) is an early study on an LSTM-based encoder for feature extraction from observation and Q-learning for action policy. LSTM-DQN++ (Yuan et al., 2018) extended the exploration and LSTM-DRQN (Yuan et al., 2018) was proposed for adding

memory units in the action scorer. KG-DQN (Ammanabrolu and Riedl, 2019) and GATA (Adhikari et al., 2020) extended the language understanding. LeDeepChef (Adolphs and Hofmann, 2020) used recurrent feature extraction along with the A2C (Mnih et al., 2016). CREST (Chaudhury et al., 2020) was proposed for pruning observation information. TWC (Murugesan et al., 2021) was proposed for utilizing common sense reasoning. However, none of these studies used the neuro-symbolic approach.

For recent neuro-symbolic RL work, the Neural Logic Machine (NLM) (Dong et al., 2018) was proposed as a method for combination of deep neural network and symbolic logic reasoning. It uses a sequence of multi-layer perceptron layers to deduct symbolic logics. Rules are combined or separated during forward propagation, and an output of the entire architecture represents complicated rules. In this paper, we compare our method with this NLM.

## 3 Proposed method

### 3.1 Problem formulation

As text-based games are sequential decision-making problems, they can naturally be applied to RL. These games are partially observable Markov decision processes (POMDP) (Kaelbling et al., 1998), where the observation text does not include the entire information of the environment. Formally, the game is a discrete-time POMDP defined by  $\langle S, A, T, R, \omega, O, \gamma \rangle$ , where  $S$  is a set of states ( $s_t \in S$ ),  $A$  is a set of actions,  $T$  is a set of transition probabilities,  $R$  is a reward function,  $\omega$  is a set of observations ( $o_t \in \omega$ ),  $O$  is a set of conditional observation probabilities, and  $\gamma$  is a discount factor. Although the state  $s_t$  contains the complete internal information, the observation  $o_t$  does not. In this paper, we follow following two assumptions: one, the word in each command is taken from a fixed vocabulary  $V$ , and two, each action command consists of two words (verb and object). The objective for the agent is to maximize the expected discounted reward  $E[\sum_t \gamma^t r_t]$ .

### 3.2 Method

The proposed method consists of two processes: converting text into first-order logic (FOL), and training the action policy in LNN.

### 3.2.1 FOL converter

The FOL converter converts a given natural observation text  $o_t$  and observation history  $(o_{t-1}, o_{t-2}, \dots)$  into first-order logic facts. The method first converts text into propositional logics  $l_{i,t}$  by a semantic parser from  $o_t$ , such as, the agent understands an opened direction from the current room. The agent then retrieves the class type  $c$  of the word meaning in propositional logic  $l_{i,t}$  by using ConceptNet (Liu and Singh, 2004) or the network of another word’s definition. For example, “east” and “west” are classified as a *direction*-type, and “coin” is as a *money*-type. The class is used for selecting the appropriate LNN for FOL training and inference.

### 3.2.2 LNN training

The LNN training component is for obtaining an action policy from the given FOL logics. LNN (Riegel et al., 2020) has logical conjunction (AND), logical disjunction (OR), and negation (NOT) nodes directly in its neural network. In our method, we prepare an AND-OR network for training arbitrary rules from given inputs. As shown in Fig. 1, we prepare all logical facts at the first layer, several AND gates (as many as the network is required) at the second layer, and one OR gate connected to all previous AND gates. During the training, the reward value is used for adding a new AND gate, and for updating the weight value for each connection. More specifically, the method is storing the replay buffer which has current observation  $o_t$ , action  $a_t$ , reward  $r_t$ , and next observation  $o_{t+1}$  value. For each training step, the method selects some replies, and it extracts first-order logical facts from current observation  $o_t$  and action  $a_t$ . The LNN trains by this fact inputs and reward; that means it forwards from input facts through LNN, calculates a loss values from the reward value, and optimizes weights in LNN. The whole training mechanism is similar to DQN (Mnih et al., 2013), the difference from these is the network. To aid the interpretability of node values, we define a threshold  $\alpha \in [\frac{1}{2}, 1]$  such that a continuous value is considered *True* if the value is in  $[\alpha, 1]$ , and *False* if it is in  $[0, 1 - \alpha]$ .

Algorithm 1 describes the whole algorithm for the proposed method.

---

### Algorithm 1 RL by FOL-LNN

---

```

1: procedure REINFORCEMENT LEARNING
2:   for  $t = 1, 2, 3, \dots$  do
3:      $o_t \leftarrow$  Observe observation
4:      $l_{t,i} \leftarrow$  Extract logic from  $o_t, o_{t-1}, \dots$ 
5:     for  $i = 1, 2, 3, \dots$  do
6:        $c \leftarrow$  Find class from ConceptNet
7:        $\theta^c \leftarrow$  Select LNN
8:        $l_{t,i}^c \leftarrow$  Convert into FOL logic
9:        $a_{t,i} \leftarrow \theta^c(l_{t,i}^c)$ 
10:    end for
11:     $a_t \leftarrow \arg \max a_{t,i}$ 
12:     $r_t, o_{t+1} \leftarrow$  Get reward and next obs
13:    Store reply  $\langle o_t, a_t, r_t, o_{t+1} \rangle$ 
14:     $\nabla \theta \leftarrow$  Update LNN from reply
15:  end for
16: end procedure

```

---

## 4 Experiments

We evaluated the proposed method on a coin-collector game in TextWorld (Côté et al., 2018) with three different difficulties (*easy*, *medium*, and *hard*). The objective of the game is to find and collect a coin which is placed in a room within connected rooms. Since we tackle a real-world game problem rather than a symbolic games, we need to extract logical facts from given natural texts for neuro-symbolic methods. We prepare the following propositional logics as extracting logical facts: *which object is found in the observation*, *which direction has already been visited*, and *which direction the agent comes from initially*. These logical values are easily calculated from visited room history and word definitions. In this experiment, we prepared 26 logical values<sup>1</sup>, and all the following neuro-symbolic methods used these value as input. For the evaluation metric, we focused on (1) the test reward value on the unseen (test) games and (2) the number of steps to achieve the goal on unseen games. Since we focus on the performance of generalization, we only use 50 small-size (level = 5) games for training, 50 unseen games from 5 different size (level = 5, 10, 15, 20, 25) games for test<sup>2</sup>, and mini-batch in training (batch size = 4). The other parameters for the game and agent follow LSTM-DQN++ (Narasimhan et al., 2015).

<sup>1</sup>26 = (5 (object) + 4 (visited) + 4 (initial))  $\times$  2 (negation)

<sup>2</sup>The agent needs to generalize the game level size

Table 1: Average reward and number of steps (**reward**: higher is better / **number of steps**: lower is better) for each epoch on 50 unseen games with three difficulty levels. These results are from moving average ( $N = 100$ ) and 5 random seeds. Training is done on only small-size games. Although neuro-only method cannot solve unseen test games, our proposed method (FOL-LNN) can solve and converge extremely faster than other SOTAs and baselines.

Epoch	Easy game			Medium game			Hard game		
	100	200	1000	100	200	2000	100	200	2000
LSTM-DQN++ *	0.07/93.4	0.10/90.9	0.12/88.6	0.00/99.9	0.00/99.6	0.03/97.3	0.00/99.9	0.00/99.9	0.04/96.6
NLM-DQN **	0.87/26.4	0.93/20.8	<b>1.00/15.0</b>	0.27/81.1	0.48/65.9	<b>1.00/29.7</b>	0.01/99.7	0.10/94.8	0.66/64.0
NN-DQN	0.91/22.8	0.95/19.0	<b>1.00/15.0</b>	0.48/65.6	0.65/54.3	<b>1.00/29.5</b>	0.19/89.0	0.28/84.0	0.97/46.3
LNN-NN-DQN	0.88/24.8	0.94/20.2	<b>1.00/15.0</b>	0.49/65.8	0.61/57.0	<b>1.00/29.6</b>	0.24/86.9	0.27/84.9	0.97/47.4
FOL-LNN (Ours)	<b>0.95/19.0</b>	<b>0.98/17.1</b>	<b>1.00/15.0</b>	<b>0.94/32.7</b>	<b>0.97/30.7</b>	<b>1.00/28.6</b>	<b>0.95/44.8</b>	<b>0.98/43.5</b>	<b>1.00/42.0</b>

\* State-of-the-art neuro-only method with a simple DQN action scorer (Narasimhan et al., 2015)

\*\* State-of-the-art neuro-symbolic method has same input as ours and other neuro-symbolic methods (Dong et al., 2018)

We prepared five methods for an evaluation of the proposed method:

- LSTM-DQN++ (Narasimhan et al., 2015): State-of-the-art neuro-only method with a simple DQN action scorer. We use this method as a baseline method for the neuro-only agent, and LSTM receives extracted embedding vector from natural text information.
- NLM-DQN (Dong et al., 2018): State-of-the-art neuro-symbolic method. The input is propositional logical values that is also used in following baselines and proposed method. The original NLM uses the REINFORCE (Williams, 1992) algorithm, but in order to handle text-based games with the same setting as the other methods, we applied the DQN algorithm. In short, the method uses an NLM layer instead of an LSTM (Hochreiter and Schmidhuber, 1997) for the encoder of the LSTM-DQN++ method. We tuned the hyper-parameters from the same search space as the original paper.
- NN-DQN: Naïve neuro-symbolic baseline method. The input of the network is propositional logical values, and it uses a multi layer perceptron as the encoder of the LSTM-DQN++.
- LNN-NN-DQN: Neuro-symbolic baseline method. The method first gets propositional logical values, it converts by LNN into some conjunction values for all combinations of given logical values, and then it inputs them

into a multi layer perceptron. It differs from NN-DQN in that LNN-NN-DQN has prepared conjunction nodes, which should lead to faster training in beginning of the training, and better interpretability after the training.

- FOL-LNN: Our neuro-symbolic method.

Table 1 shows the test reward and test step values on unseen games, and Fig. 2 shows curves. First, all the RL results with logical input were better than those with textual input. Second, our proposed method could converge much faster than the other neuro-symbolic state-of-the-art and baseline methods. Third, only our method could extract the trained rules by checking the weight value of the LNN. We attached the extracted rules from the *medium* level games here:

$$\exists x \in W_{money} \quad \langle \text{find } x \rangle \rightarrow \langle \langle \text{take } x \rangle \rangle,$$

$$\begin{aligned} \exists x \in W_{direction} \\ & (\langle \text{find } x \rangle \wedge \neg \langle \text{visited } x \rangle \wedge \neg \langle \text{initial } x \rangle) \vee \\ & (\langle \text{find } x \rangle \wedge \langle \text{all are visited} \rangle \wedge \langle \text{initial } x \rangle) \rightarrow \langle \langle \text{go } x \rangle \rangle, \end{aligned}$$

where  $W_{direction}$  is a set of words in a type of “direction” in ConceptNet. The rule for “take”-action is for taking a coin. The first conjunction rule for “go”-action is for visiting an un-visited room, and the second rule is for returning to the initial room from a dead-end. With our proposed method, we can see that these trained rules will be helpful for operating the neural agent in real use cases.

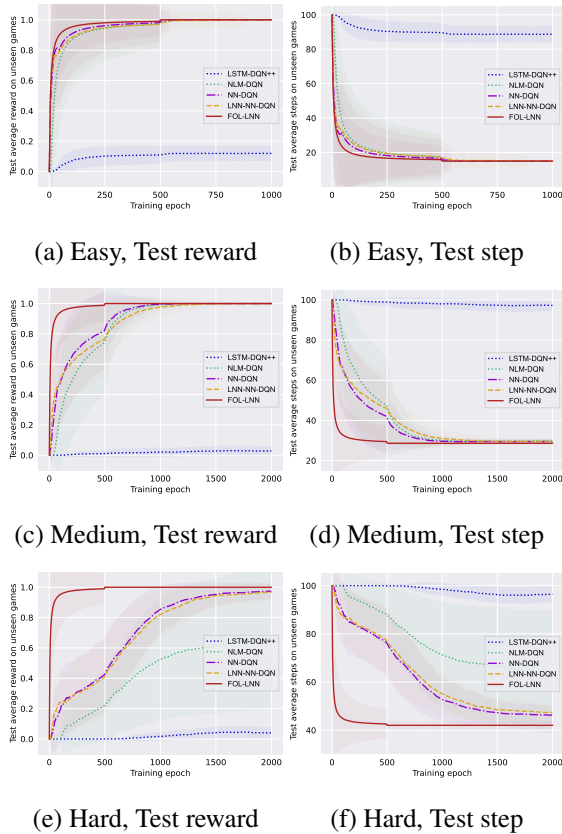


Figure 2: Curves for reward and number of steps for 50 unseen games. Moving average is applied.

## 5 Conclusion

In this paper, we proposed a novel neuro-symbolic method for RL on text-based games. According to the evaluation on the natural language text-based game with several difficulties, our method can converge extremely faster than other state-of-the-art neuro-only and neuro-symbolic methods, and extract trained logical rules for improving interpretability of the model.

## Discussion about ethics

Our model is not using any sensitive contexts such as legal or health-care settings. The data set used in our experiment does not contain any sensitive information. Since our proposed neuro-symbolic RL method can extract the trained rules for interpretability of the model, the method can analyze a reason behind taken action. We are sure that if the model returns biased results, this functionality is helpful for clearing the reason for these data bias issues.

## References

- Ashutosh Adhikari, Xingdi Yuan, Marc-Alexandre Côté, Mikuláš Zelinka, Marc-Antoine Rondeau, Roman Laroche, Pascal Poupart, Jian Tang, Adam Trischler, and William L. Hamilton. 2020. [Learning dynamic belief graphs to generalize on text-based games](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Leonard Adolphs and Thomas Hofmann. 2020. Ledeechef deep reinforcement learning agent for families of text-based games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 7342–7349.
- Prithviraj Ammanabrolu and Mark Riedl. 2019. [Playing text-adventure games with graph-based deep reinforcement learning](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 3557–3565.
- Subhajit Chaudhury, Daiki Kimura, Kartik Talamadupula, Michiaki Tatsubori, Asim Munawar, and Ryuki Tachibana. 2020. [Bootstrapped q-learning with context relevant observation pruning to generalize in text-based games](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 3002–3008.
- Marc-Alexandre Côté, Ákos Kádár, Xingdi Yuan, Ben Kybartas, Tavian Barnes, Emery Fine, James Moore, Matthew J. Hausknecht, Layla El Asri, Mahmoud Adada, Wendy Tay, and Adam Trischler. 2018. [Textworld: A learning environment for text-based games](#). In *Computer Games - 7th Workshop, CGW 2018, Held in Conjunction with the 27th International Conference on Artificial Intelligence, IJCAI 2018, Stockholm, Sweden, July 13, 2018, Revised Selected Papers*, pages 41–75.
- Honghua Dong, Jiayuan Mao, Tian Lin, Chong Wang, Lihong Li, and Denny Zhou. 2018. Neural logic machines. In *International Conference on Learning Representations*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. 1998. [Planning and acting in partially observable stochastic domains](#). *Artificial Intelligence*, 101(1):99–134.
- Daiki Kimura. 2018. Daqn: Deep auto-encoder and q-network. *arXiv preprint arXiv:1806.00630*.
- Daiki Kimura, Subhajit Chaudhury, Ryuki Tachibana, and Sakyasingha Dasgupta. 2018. [Internal model from observations for reward shaping](#).

- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- H. Liu and P. Singh. 2004. [Conceptnet — a practical commonsense reasoning tool-kit](#). *BT Technology Journal*, 22(4):211–226.
- Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. 2016. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937. PMLR.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin A. Riedmiller, Andreas Fiedjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dhharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. 2015. Human-level control through deep reinforcement learning. *Nature*, 518:529–533.
- Keerthiram Murugesan, Mattia Atzeni, Pavan Kapani-pathi, Pushkar Shukla, Sadhana Kumaravel, Gerald Tesauero, Kartik Talamadupula, Mrinmaya Sachan, and Murray Campbell. 2021. Text-based rl agents with commonsense knowledge: New challenges, environments and baselines. In *Thirty Fifth AAAI Conference on Artificial Intelligence*.
- Karthik Narasimhan, Tejas D. Kulkarni, and Regina Barzilay. 2015. [Language understanding for text-based games using deep reinforcement learning](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 1–11.
- Ryan Riegel, Alexander Gray, Francois Luus, Naweed Khan, Ndivhuwo Makondo, Ismail Yunus Akhalwaya, Haifeng Qian, Ronald Fagin, Francisco Barahona, Udit Sharma, Shajith Iqbal, Hima Karanam, Sumit Neelam, Ankita Likhyani, and Santosh Srivastava. 2020. [Logical neural networks](#).
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.
- Xingdi Yuan, Marc-Alexandre Côté, Alessandro Sordani, Romain Laroche, Remi Tachet des Combes, Matthew J. Hausknecht, and Adam Trischler. 2018. [Counting to explore and generalize in text-based games](#). *CoRR*, abs/1806.11525.

## Appendix A: Environment Setting

In this section we describe Coin-Collector (Yuan et al., 2018), a text-based game used in our experiments. Then, we describe the hyper-parameter setting.

### Appendix A.1: Coin-Collector

Coin-Collector is a kind of text-based games, and we have to move an agent through rooms to get a coin placed in a room. An agent receives an observation text that describes the structure of a current room from a game. The goal of Coin-Collector is to analyze textual data and understand the structure of given rooms for training an agent.

A game has hyper-parameters such as *level* and *difficulty*. A game *level* indicates the minimum number of steps to a room in which a coin is placed. Rooms are randomly connected and their structure depends on *difficulty*. An easy game has no distractor rooms (dead ends) along the path. On a medium game, each room along the optimal trajectory has one distractor room randomly connected to it. A hard game, each room has two distractor rooms which means each room has one for optimal trajectory, one for the previous room, and two for distractor rooms.

An agent can use two types of verbs ({take, go}) and five types of nouns ({coin, east, west, south, north}). Since an action consists of a verb and a noun, there are ten different actions that an agent can take. For the settings of LSTM-DQN++ (Narasimhan et al., 2015), the agent gets the positive reward when the agent goes in a new room. The agent also gets positive reward when the agent successfully returns the initial coming direction for medium setting. If an agent takes an invalid action such as “go coin”, or “go north” at no north room, the agent does not receive a negative reward.

### Appendix A.2: Hyper-parameters

For the all experiments, we used the same hyper-parameters with the previous work for Coin-Collector as follows.

- We used a prioritized replay memory with capacity of 500,000 and the priority fraction is 0.25.
- A mini-batch gradient update is performed every 4 steps in the game play.
- The discount factor for Q-learning  $\gamma$  is 0.9.

- We used an episodic discovery bonus that encourages an agent to discover unseen states and the coefficient  $\beta$  is 1.0.
- We anneal the  $\epsilon$  for the  $\epsilon$ -greedy strategy from 1 to 0.2 over 1000 epochs. After 1000 episodes, the  $\epsilon$  is 0.2.
- We used the Adam algorithm (Kingma and Ba, 2014) for the optimization and the learning rate is  $1e^{-3}$ .

## Appendix B: Experiment details

The training and validation times until 3,000 epochs for each method are as follows.

- LSTM-DQN++ (Narasimhan et al., 2015): Around 2 hours for easy difficulty, and around 4 hours for medium difficulty.
- NLM-DQN (Dong et al., 2018): Around 40 minutes for easy difficulty, and around 2.5 hours for medium difficulty.
- NN-DQN: Around 30 minutes for easy difficulty, and around 1.5 hours for medium difficulty.
- LNN-NN-DQN: Around 30 minutes for easy difficulty, and around 1.5 hours for medium difficulty.
- FOL-LNN: Around 35 minutes for easy difficulty, and around 2 hours for medium difficulty.

These results are calculated on Intel Core i7-6700K CPU (4.00GHz) and NVIDIA Titan X. From these results, our proposed method is not computationally expensive than other methods.