Neuro-Symbolic Agent with ASP for Robust Exception Learning in Text-Based Games

Kinjal Basu[1]

[1]IBM Research, USA

Abstract

Text-based games (TBGs) present a significant challenge in natural language processing (NLP) by requiring reinforcement learning (RL) agents to combine language comprehension with reasoning. A primary difficulty for these agents is achieving generalization across multiple games, particularly in handling both familiar and novel objects. While deep RL approaches excel in scenarios with known objects, they struggle with unseen ones. Commonsense-augmented RL agents address this but often lack interpretability and transferability. To address these limitations, we propose a neuro-symbolic framework that integrates symbolic reasoning with neural RL models. Our approach utilizes inductive logic programming (ILP) to learn symbolic rules dynamically. We show that this hybrid agent outperforms pure neural agents in handling familiar objects. Additionally, we introduce a novel generalization approach based on information gain and WordNet that helps our agent to excel in the test sets with unseen objects as well.

Keywords

Answer Set Programming, Reinforcement Learning, Text-Based Games

1. Introduction

Natural language is key to human intelligence, and text-based games (TBGs) offer a platform to study language-driven decision-making in AI. To succeed in TBGs, agents need both natural language processing (NLP) and reinforcement learning (RL) skills. Existing RL agents are either rule-based, relying on predefined knowledge, or neural-based, using deep learning [1]. Rule-based agents lack flexibility, while neural agents require large training sets, learn slowly, and struggle to generalize to unseen entities. Additionally, neural RL agents' policies are not interpretable.

In this paper (it is an extended abstract of the EXPLORER paper [2]), we introduce a neuro-symbolic architecture for text-based games (TBGs) that combines the strengths of both neural and symbolic agents. The symbolic agent learns interpretable policies from game interactions, allowing for non-monotonic reasoning (NMR), where the agent's beliefs adapt based on new information. The neural component explores the environment and assists when symbolic rules are lacking. The neuro-symbolic agent demonstrates strong performance on the TextWorld cooking domain dataset but struggles on the TextWorld Commonsense (TWC) framework due to out-of-distribution (OOD) challenges. To address this, the symbolic rules are generalized

4th Workshop on Goal-directed Execution of Answer Set Programs (GDE'24), October 12, 2024
Kinjal.Basu@ibm.com (K. Basu)
0000-0001-8693-9307 (K. Basu)

using WordNet, enabling the agent to handle unseen objects in the OOD test set, resulting in superior performance compared to state-of-the-art methods on TWC.

2. Methodology

This paper presents a hybrid approach to reinforcement learning (RL) in text-based games (TBGs) by combining neural and symbolic agents. Neural agents excel in exploration, while symbolic agents specialize in learning interpretable policies based on rewards. By leveraging both, the system aims to achieve more effective results. The symbolic agent learns logic-based policies and applies them via an Answer Set Programming (ASP) solver, with the neural agent stepping in when the symbolic agent cannot provide suitable actions. A key goal of the system is to ensure robust performance on both seen and unseen (out-of-distribution) entities. To address this, the paper introduces a novel approach for policy generalization, dynamically generating generalized rules using WordNet hypernym relations.

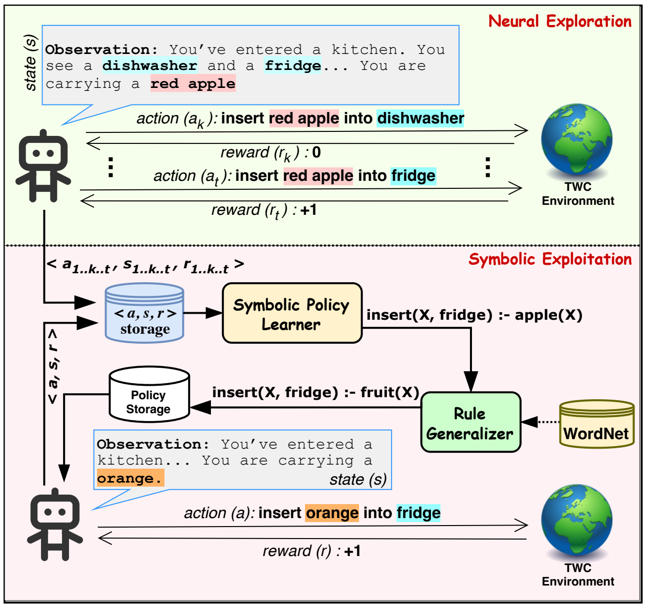

Figure 1: NeuroSymbolic Agent

3. Experiments and Results

Data and Experimental Setup: To test our neuro-symbolic agent, we choose TW-Cooking domain [3], which requires both exploration and exploitation. As the name suggests, this game suit is about collecting various cooking ingredients and preparing a meal following an in-game recipe. To showcase the generalization capability, we have tested our *neuro-symbolic* and *neuro-symbolic + generalization* agents on TWC games with OOD data. With the help of TWC framework [4], we have generated a set of games with 3 different difficulty levels: *easy*, *medium*, and *hard*.

Results: We have tested the cooking game with *Neuro-Symbolic Rules* and compared the results with the baseline model (LSTM-A2C). Table 1 illustrates the results, where 'L' means *difficulty level*. For TWC, table 2 shows the comparison results of all the 4 settings along with the baseline model (text-only agent). We compared our agents in two different test sets - (i) *IN distribution:* that has the same entities as the training dataset, and (ii) *OUT distribution:* that has new entities, which have not been included in the training set.

Table 1
Comparison Results for the TW-Cooking Domain

Level	Text-Only (Neural)		Neural + Symbolic Rules (Neuro-Symbolic)	
	# Steps	N. Score	# Steps	N. Score
L-1	11.50 ± 1.26	0.93 ± 0.01	$\mathbf{9.17 \pm 3.28}$	$\mathbf{0.96 \pm 0.04}$
L-2	45.79 ± 1.56	0.35 ± 0.07	$\mathbf{15.60 \pm 3.74}$	$\mathbf{0.58 \pm 0.04}$
L-3	46.91 ± 1.51	0.25 ± 0.01	$\mathbf{26.92 \pm 2.74}$	$\mathbf{0.34 \pm 0.01}$
L-4	22.03 ± 0.19	0.76 ± 0.01	$\mathbf{19.85 \pm 2.12}$	$\mathbf{0.85 \pm 0.07}$

Table 2
Results for TWC - Within Distribution (IN) and Out-of-Distribution (OUT) Games. We report the number of steps taken by the agent (lower is better) and the normalized scores (higher is better).

	Method	Easy		Medium		Hard	
		Steps	N. Score	Steps	N. Score	Steps	N. Score
IN	Text Only	15.12 ± 1.95	0.91 ± 0.03	33.17 ± 2.76	0.83 ± 0.04	47.68 ± 2.43	0.60 ± 0.05
	Neural + Symbolic Rules	17.39 ± 3.01	0.93 ± 0.04	46.7 ± 2.14	0.42 ± 0.12	37.66 ± 0.93	0.88 ± 0.01
	Neural + Generalized Rules (Exhaustive)	12.86 ± 3.04	0.91 ± 0.04	29.9 ± 3.16	0.65 ± 0.06	**30.44 ± 0.87**	**0.95 ± 0.03**
	Neural + Generalized Rules (IG Hyp. Lvl. 2)	10.59 ± 1.3	0.95 ± 0.02	**22.57 ± 1.04**	**0.77 ± 0.07**	30.46 ± 0.74	0.87 ± 0.01
	Neural + Generalized Rules (IG Hyp. Lvl. 3)	**9.55 ± 2.34**	**0.96 ± 0.02**	25.34 ± 2.86	0.76 ± 0.03	33.54 ± 1.47	0.91 ± 0.03
OUT	Text Only	16.66 ± 1.74	0.92 ± 0.03	37.3 ± 3.45	**0.73 ± 0.06**	50.00 ± 0.00	0.30 ± 0.04
	Neural + Symbolic Rules	21.19 ± 0.87	0.84 ± 0.06	46.36 ± 1.52	0.42 ± 0.08	44.25 ± 0.42	0.63 ± 0.01
	Neural + Generalized Rules (Exhaustive)	14.65 ± 2.18	0.91 ± 0.05	**37.07 ± 2.09**	0.63 ± 0.06	**41.52 ± 1.12**	**0.83 ± 0.02**
	Neural + Generalized Rules (IG Hyp. Lvl. 2)	15.08 ± 1.2	0.91 ± 0.02	40.63 ± 3.03	0.57 ± 0.06	42.18 ± 0.66	0.79 ± 0.01
	Neural + Generalized Rules (IG Hyp. Lvl. 3)	**12.72 ± 1.22**	**0.92 ± 0.02**	37.38 ± 3.09	0.64 ± 0.09	43.16 ± 2.83	0.78 ± 0.03

4. Conclusion

This paper presents a neuro-symbolic approach that integrates a non-monotonic reasoning (NMR) symbolic agent with a neural agent in a text-based reinforcement learning (RL) environment. We introduce a novel approach for rule generalization based on information gain. Our method not only yields encouraging results in the TW-Cooking domain and TWC games but also produces interpretable and transferable policies.

References

[1] K. Narasimhan, T. Kulkarni, R. Barzilay, Language understanding for text-based games using deep reinforcement learning, arXiv preprint arXiv:1506.08941 (2015).

[2] K. Basu, K. Murugesan, S. Chaudhury, M. Campbell, K. Talamadupula, T. Klinger, Explorer: Exploration-guided reasoning for textual reinforcement learning, arXiv preprint arXiv:2403.10692 (2024).

[3] A. Adhikari, X. Yuan, M.-A. Côté, M. Zelinka, M.-A. Rondeau, R. Laroche, P. Poupart, J. Tang, A. Trischler, W. Hamilton, Learning dynamic belief graphs to generalize on text-based games, Advances in Neural Information Processing Systems 33 (2020) 3045–3057.

[4] K. Murugesan, M. Atzeni, P. Kapanipathi, P. Shukla, S. Kumaravel, G. Tesauro, K. Talamadupula, M. Sachan, M. Campbell, Text-based rl agents with commonsense knowledge: New challenges, environments and baselines, in: Thirty Fifth AAAI Conference on Artificial Intelligence, 2021.