

Creativity of AI: Automatic Symbolic Option Discovery for Facilitating Deep Reinforcement Learning

Mu Jin¹, Zhihao Ma¹, Keping Jin¹, Hankz Hankui Zhuo^{1*}, Chen Chen², Chao Yu¹

¹ School of Computer Science and Engineering, Sun Yat-Sen University

² Huawei Noah's Ark Lab

{jinm6, mazhh7, jinkb}@mail2.sysu.edu.cn, {zhuohank, yuchao3}@mail.sysu.edu.cn, chenchen9@huawei.com

Abstract

Despite of achieving great success in real life, Deep Reinforcement Learning (DRL) is still suffering from three critical issues, which are data efficiency, lack of the interpretability and transferability. Recent research shows that embedding symbolic knowledge into DRL is promising in addressing those challenges. Inspired by this, we introduce a novel deep reinforcement learning framework with symbolic options. This framework features a loop training procedure, which enables guiding the improvement of policy by planning with action models and symbolic options learned from interactive trajectories automatically. The learned symbolic options alleviate the dense requirement of expert domain knowledge and provide inherent interpretability of policies. Moreover, the transferability and data efficiency can be further improved by planning with the action models. To validate the effectiveness of this framework, we conduct experiments on two domains, Montezuma's Revenge and Office World, respectively. The results demonstrate the comparable performance, improved data efficiency, interpretability and transferability.

Introduction

Deep Reinforcement Learning (DRL) has achieved tremendous success in complex and high dimensional environments such as Go (Silver et al. 2016, 2017) and Atari Games (Mnih et al. 2015). It interacts with environments and improves its policy with the collected experience, by maximizing the long term reward. Recent criticism on DRL mostly focuses on the lack of transferability, interpretability, and data efficiency. The policy learnt from an environment often fails in another unseen environment. Due to the use of black-box neural networks for function approximation, the intrinsic lack of interpretability issue naturally raises in DRL, which disables the agent to explain its actions in a human-understandable way and earn people's trust in critical areas such as autonomous driving (Aradi 2020) and chemical engineering (Zhou, Li, and Zare 2017). Besides, DRL often requires a large amount of data to learn a satisfying policy in complex environments. The process of collecting experiences for learning is time-consuming and the sample efficiency is low.

*Corresponding author

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: State Mapping Function in Montezuma's Revenge

To alleviate those issues, researchers have investigated the combination of HRL and symbolic planning to improve transferability, interpretability, and data efficiency (Ryan 2002; Leonetti, Iocchi, and Stone 2016; Yang et al. 2018; Lyu et al. 2019; Illanes et al. 2020; Sarathy et al. 2020; Lee et al. 2021). In those works, the original MDP is divided into two levels. The higher level utilizes a symbolic planner with a given action model to generate plans for selecting options, while the lower level interacts with the environment to accomplish the selected options (Illanes et al. 2020; Lee et al. 2021). The two-level structure helps alleviating the sparse reward issue, and improves sampling efficiency with the help of generated plans. In those works, they all assume the action models have been provided by domain experts. In many real-world applications, however, it is often difficult to create action models by hand, especially when the environment is complicated. A more realistic idea is to automatically learn action models from training data (Zhuo and Kambhampati 2017; Yang, Wu, and Jiang 2007; Ng and Petrick 2019; Martínez et al. 2016; James, Rosman, and Konidaris 2020) and exploit the learnt action models to generate plans for guiding the exploration of options. Although there is indeed an approach (Sarathy et al. 2020) proposed to learn action models automatically, they still need to manually define major parts of action models in advance. Besides, the planning goal in this approach is kept unchanged while in our framework it is dynamically adapted to maximize the external reward.

In this work, we propose a novel framework, namely SORL, which stands for Symbolic Options for Reinforcement Learning, to learn action models to help the exploration of actions in reinforcement learning. We assume that there exists

<p>action act0: pre^+: (atMiddleLadder)(keyexist) pre^-: \emptyset eff^+: (atRightLadder) (increase (reward) 0) eff^-: (atMiddleLadder)</p>
--

Figure 2: An action model in Montezuma’s Revenge

a function F , mapping high dimensional states to symbolic states and enabling us to learn symbolic action models and options. As shown in Figure 1, we extract the position of the man in red and the key from the high-dimensional state to obtain the corresponding symbolic state. When the agent walks from the middle ladder to the right ladder, the key still exists and the environment does not give any feedback (e.g. zero reward). This can be seen as a symbolic transition and we will generate the corresponding action model as shown in Figure 2. Then, we use a planner with the learned action models and planning goals as input to generate plans and use them to instruct the learning of the agent.

Based on the assumption, SORL features a two-level structure, of which the higher level is a symbolic planner and a meta-controller, and the lower level is an RL agent interacting with the environment. The higher level utilizes the collected trajectories from the lower level to learn action models and symbolic options with minimum human knowledge. After that, the meta-controller chooses an option according to the plan generated from the planner with the learned action models and assigns it to the lower level. By interacting with the environments, the lower level learns a policy to reach the assigned option and sends the collected experience to the higher level. This cross-fertilization structure not only helps alleviating the sparse and delayed reward problem but also improves the data efficiency.

Different from previous action model learning approaches (Zhuo and Kambhampati 2017; Yang, Wu, and Jiang 2007), which assume the number of action models to be learnt is known beforehand, our SORL does not know exactly “how many” and “what” action models to be learnt from the environment. We expect the agent continuously creates new action models via the interactions with the environment, and exploits both new and old action models to guide the exploration of actions to build policies creatively—we consider this as a *creativity* property of an AI agent. It was previously mentioned that AI techniques can be used to create new ideas in three ways: producing novel combinations of familiar ideas (e.g., poetic imagery and analogy), exploring the potential of conceptual spaces (to generate new ideas), and making transformations (of structured conceptual spaces) for enabling the generation of previously impossible ideas (Boden 1998). Current AI techniques have exhibited great progress in exploring the potential of conceptual spaces (the second way) by considering it as combinational heuristic search, provided the domain model or knowledge has been built by domain experts (Boden 2009). There are, however, very few works that investigate the creativity that is capable of automatically accumulating (or learning) “pieces” of knowledge (e.g., ac-

tion models) from environments and improving the ability of solving real-world application problems with the accumulated knowledge. In this work, we claim that an agent with such creativity can build better policies with respect to transferability, interpretability, and data-efficiency.

We summarize our contribution as below:

- Our work is the first one to learn action and option models automatically without being told any knowledge of these models and simultaneously learn RL policies.
- We propose a symbolic reinforcement learning framework capable of providing transferability, interpretability, and improved data-efficiency.
- The symbolic option learned by SORL is more general, which can correspond to more than one action model.

Preliminaries

In this chapter, we establish relevant notation and briefly introduce key aspects of symbolic planning and reinforcement learning.

Symbolic Planning with PDDL

In PDDL language, states are represented as set of propositions and we call it symbolic states throughout the paper to distinguish them from states in RL. Propositions represent the properties of the world and in the symbolic state s , proposition $p \in s$ if p is true otherwise ($not\ p$) $\in s$. An action description called action model is a tuple $(name, pre^+, pre^-, eff^+, eff^-)$, where $name$ is the name of the action, (pre^+, pre^-) are the preconditions and (eff^+, eff^-) are the effects. As shown in Fig.2, the action model describes that when the agent walks from the middle ladder to the right ladder, the key keeps still and the reward remains unchanged. If $pre^+ \subset s$ and $s \cap pre^- = \emptyset$, then we can execute action a and obtain the next state

$$s' = ((s - eff^-) \cup eff^+).$$

The planning domain $D = (P, A)$ includes the proposition set P and the action set A , which describe the state space and the action space, respectively. A tuple (s, a, s') describes a symbolic transition from state s to state s' after executing action a . We define a planning problem denoted as a triple (I, P, A, G) , of which I is an initial state and G is a goal state. The solution to this problem is called a plan π , which is a sequence of actions. After executing the plan, we can obtain a symbolic transition trace from I to G . To obtain such a plan with the maximum reward, we use a planner called Metric-FF (Hoffmann 2002), which can handle planning problems with continuous metrics.

Reinforcement Learning

A Markov Decision Process (MDP) is defined as the tuple $(\tilde{S}, \tilde{A}, P_{\tilde{s}\tilde{s}'}^{\tilde{a}}, r_{\tilde{s}}^{\tilde{a}}, \gamma)$ where \tilde{S} and \tilde{A} denote the state space and action space, respectively, $P_{\tilde{s}\tilde{s}'}^{\tilde{a}}$ provides the transition probability of moving from state $\tilde{s} \in \tilde{S}$ to state $\tilde{s}' \in \tilde{S}$ after taking action $\tilde{a} \in \tilde{A}$, $r_{\tilde{s}}^{\tilde{a}}$ is the immediate reward obtained after performing action \tilde{a} at state \tilde{s} and $\gamma \in [0, 1)$ is a discount factor.

The task of RL is to obtain a policy

$$\pi : \tilde{S} \rightarrow \tilde{A}$$

that maximizes the expected return

$$V_\pi(\tilde{s}) = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t r_t \mid \tilde{s}_0 = \tilde{s} \right]$$

where r_t is the reward at time step t received by following π from state $\tilde{s}_0 = \tilde{s}$. The state-action value function is defined as follows:

$$Q_\pi(\tilde{s}, \tilde{a}) = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t r_t \mid \tilde{s}_0 = \tilde{s}, \tilde{a}_0 = \tilde{a} \right].$$

Option Framework

Hierarchical Reinforcement Learning (HRL) extends RL with temporally macro actions that represent high-level behaviors. The option framework (Sutton, Precup, and Singh 1999) models macro actions as options. In particular, an option o is defined as $(I_o(s), \pi_o(s), \beta_o(s))$, where initiation condition $I_o(s)$ determines whether the option o can be executed at state s , termination condition $\beta_o(s)$ determines whether option execution terminates at state s and $\pi_o(s)$ is a policy mapping state s to a low-level action. In this framework, an agent learns to choose an optimal option to be executed in the higher level, i.e. meta controller level and the lower level, i.e. controller level learns optimal policies to reach the option. An explicit assumption is that the set of options is predefined by human experts.

The SORL Framework

We define the reinforcement learning environment by a tuple $(I, G, P, A, F, \tilde{S}, \tilde{A}, \tilde{R}, \tilde{P}, \gamma)$ and divide it into three parts:

- First, we define a high-level symbolic planning problem by (I, G, P, A) , where I is an initial state, and G is a goal state. P is a set of propositions represented by planning language PDDL with prior knowledge and it is used to describe symbolic states S , where $S \subseteq 2^P$. A is a set of action models that $S \times A \rightarrow S$ transfer a symbolic state to another. Each action model is learned by meta-controller through symbolic state pairs.
- Second, we define a state mapping function $F : \tilde{S} \rightarrow 2^P$ mapping a high dimension state \tilde{s} to a symbolic state s .
- At last, we define an underlying decision-making problem by an MDP tuple $(\tilde{S}, \tilde{A}, \tilde{R}, \tilde{P}, \gamma)$. We denote a symbolic action and state as a and s respectively, while the primitive action and state as \tilde{a} and \tilde{s} . Noted that \tilde{a} and \tilde{s} are gained from interacting with the environment.

Taking the Figure 1 as an example, in the game of the Montezuma's Revenge, the high level dimension state is the picture of the game scene and the symbolic state is composed of the propositions describing the location of the agent and the existence of the key.

This framework aims to learn action models, which can be utilized by logic-based planner to generate a sequence of options and achieve the maximal cumulative reward. As shown

in the Figure 3, the SORL framework includes three components: (1) a planner for generating plans, (2) a meta-controller for generating action models, goals and choosing the goal option, and (3) an option set for interacting with the environment. The meta-controller first takes the symbolic state pairs and their external rewards as inputs and outputs action models and a goal. Noted that the state pairs set are empty in the beginning. Then the planner takes the action models as input and computes a plan. Next, the meta-controller receives the plan from the planner and chooses an option. Each option in option set can be regarded as an agent. The chosen agent keeps interacting with the environment until accomplishing the option or reaching the maximal steps, and the low-level state traces will be transformed into symbolic state pairs by the label function F and sent back to meta controller. The meta-controller continues learning action models and symbolic options from gained symbolic state pairs and external rewards. We repeat these procedures $num_episodes$ times. With the proceeding of learning, our approach keeps updating action models and planning goals and the planner is able to generate plans achieving better rewards.

Option Set

Symbolic Option In this paper, we propose a novel option framework which is called *symbolic option*. A symbolic option is computed by symbolic state pairs gained from trajectories instead of manual setting in advance, requiring less prior knowledge in our approach. We define a symbolic option by $so = (pre, \pi, eff)$, where π is a low-level policy and pre is an union of preconditions, including pre^+ and pre^- . It is created and updated when the meta-controller generates action models. Similarly, eff is composed of eff^+ and eff^- , describing the effects of the symbolic option. As for a symbolic option so and a high-dimension state \tilde{s} , we compute initiation condition $I_{so}(\tilde{s})$ by Equation (1) and termination condition $\beta_{so}(\tilde{s})$ by Equation (2). A symbolic option can be executed based on \tilde{s} only if $I_{so}(\tilde{s}) = True$. Similarly, it terminates only if $\beta_{so}(\tilde{s}) = True$.

$$I_{so}(\tilde{s}) = \begin{cases} True & pre^+ \subset F(\tilde{s}), F(\tilde{s}) \cap pre^- = \emptyset \\ False & otherwise \end{cases} \quad (1)$$

$$\beta_{so}(\tilde{s}) = \begin{cases} True & eff^+ \subset F(\tilde{s}), eff^- \cap F(\tilde{s}) = \emptyset \\ False & otherwise \end{cases} \quad (2)$$

It is noted that the inherent symbolic propositions of our symbolic option provide better interpretability compared to those approaches based on black-box neural networks. In terms of the low-level policy π , it can be learned by interacting with the environments with the intrinsic rewards given by the meta-controller.

Global Option At the beginning of our algorithm, the option set contains no symbolic options but a global option

$$o_G = (I_G(s), \pi, \beta_G(s)),$$

where $I_G(\tilde{s}) \equiv True$, $\beta_G(\tilde{s}) = True$ if symbolic state changes and $\pi = random(\tilde{A})$. We use $random(\tilde{A})$ to indicate that the global option each time chooses a random




Figure 3: The SORL framework

action $\tilde{a} \in \tilde{A}$. Intuitively, the global option is available for any state and it keeps randomly exploring until the symbolic state changes. Hence, in order to discover new action models, the meta-controller outputs the global option when the plan is empty or all action models in the plan has been executed.

Symbolic State Pair and External Reward Given an option o_j under state \tilde{s} , the lower level policy interacts with the environment and output a pair of symbolic states (s_1, s_2) and external reward r_e , denoted by

$$(s_1, s_2), r_e = \text{ExecuteOption}(\tilde{s}, o_j).$$

If the chosen option o_j isn't available for \tilde{s} , i.e., $I_j(\tilde{s}) = \text{False}$, both of the output pair and the reward are *None*. Otherwise, if the chosen option o_j is able to be executed, we let $s_1 = F(\tilde{s})$ and the policy π_j first chooses an action \tilde{a} and we can obtain the next state \tilde{s}' and its reward \tilde{r} by interacting with the environments. Then the controller adds experience $(\tilde{s}, \tilde{a}, \tilde{s}', \tilde{r})$ to the o_j 's replay buffer. We keep executing action by following the low-level policy and update the states and rewards until $\beta_j(\tilde{s}') = \text{True}$ or reaching the maximum steps, which means the option has been successfully executed or not. Finally, if the option o_j is successfully executed, we set the output symbolic state pair as (s_1, s_2) of which $s_2 = F(\tilde{s}')$ and the external reward r_e be the accumulated sum of the environment rewards during interacting.

Meta-controller

In this section, we introduce our Meta-controller in detail. Meta-controller takes symbolic state pairs and their external rewards as input, and first generates action models and a planning goal and then chooses a an option according to the plan from planner.

Action Model Given symbolic state pairs and their rewards, meta-controller generates action models by

$$A, F_{A,O}, O = \text{GenerateActionModels}(R, O, sr).$$

The function indicates it takes a dictionary R , an option set O and the success ratio set sr as inputs, and outputs a generated action set A , a mapping function $F_{A,O}$ and the updated option set O . Dictionary R includes mappings from a symbolic state pair to its external rewards. $F_{A,O}$ transfers action models to options. The success ratio set sr records the percentage of action models successful executed each 100 times.

As for a symbolic state pair $(s_1, s_2)_i \in R$, we can get a corresponding action model

$$a_i = (\text{name}, pre^+, pre^-, eff^+, eff^-).$$

Note that the action model and a symbolic state pair are a one-to-one match. Given a state pair $(s_1, s_2)_i$, the *name* of a_i is the index of action models, denoted by act_i , and $pre^+ = \{p|p \in s_1\}$, $pre^- = \{p|p \notin s_1\}$. Let $eff^+ = s_2 - s_1$ and $eff^- = s_1 - s_2$, where $a - b$ is a set subtraction indicating set a subtracts the intersection of set a and set b . In order to generate a plan gaining a maximum reward, we use the metric constant *quality* to denote the cumulative reward of the plan and add the proposition “(increase (quality) ρ_i)” into eff^+ . Finally, we get an action which is called act_i , and we define the gained reward of act_i by ρ_i . To encourage the planner to generate a plan including the exploring action model, the reward ρ_i is composed of mean external reward and exploration reward, computed by Equation (3), where $R[(s_1, s_2)_i]$ is the external rewards list and r_E is the exploration rewards. The exploration rewards is computed by Equation (4), where c is a constant and $sr[i]$ is the success rate of act_i , which means the exploration reward decreases as the success rate increases.

$$\rho_i = \text{mean}(R[(s_1, s_2)_i]) + r_E \quad (3)$$

$$r_E = \begin{cases} c(1 - sr[i]) & act_i \text{ is being explored} \\ 0 & otherwise \end{cases} \quad (4)$$

If there exists a symbolic option $o_j = (pre_j, \pi_j, eff_j)$ where $eff_j = eff$ after we attain an action model, we update pre_j^+ to a union of pre^+ and pre_j^+ , and pre_j^- to a union of pre^- and pre_j^- . Otherwise, we create a new symbolic option $o_j = (pre, \pi_j, eff)$ and add it to the option set O . Finally, we set the mapping function $F_{A,O}(act_i) = o_j$. During the exploration, we explore each action model sequentially, in other words, we repeat exploring act_i until the success rate of act_0 to act_{i-1} is higher than the threshold.

Planning Goal Next Meta-controller outputs a goal to guide planner, aiming at generating a plan with a maximal reward. The goal is a label function *quality* $> q$, where q is the cumulative external rewards of the plan gained in the last episode. Intuitively, the function constrains the planner to compute a plan with a largest reward compared with the past plans.

Chosen Option and Intrinsic Reward After the planner generates a plan $\Pi = (a_1, a_2, \dots, a_n)$, as for each action model a_i , the meta-controller selects a symbolic option from option set by $o_j = F_{A,O}(act_i)$, and we can get an series of

Algorithm 1: Planning and Learning algorithm for SORL

Input: proposition set P , state mapping function F , success ratio threshold λ

```
1: Initialization: option set  $O \leftarrow \{o_G\}$ , action models set  
    $A \leftarrow \emptyset$ , symbolic state pairs' external rewards dictionary  
    $R \leftarrow \emptyset$ , action models' success ratio set  $sr \leftarrow \emptyset$ , plan  
    $\Pi_0 \leftarrow \emptyset$ ,  $q \leftarrow 0$   
2: for  $t=1,2, \dots, num\_episodes$  do  
3:   Initialize game, get start state  $\tilde{s}_0$ ,  $I \leftarrow F(\tilde{s}_0)$ ,  $\Pi^* \leftarrow$   
    $\Pi_{t-1}$   
4:    $A, F_{A,O}, O \leftarrow GenerateActionModels(R, O, sr)$   
5:    $G \leftarrow (quality > q)$   
6:    $\Pi_t \leftarrow metricFF.solve(I, P, A, G)$   
7:   if  $\Pi_t = \emptyset$  then  $\Pi_t \leftarrow \Pi^*$   
8:    $q \leftarrow 0$   
9:   for  $a_i \in \Pi_t$  do  
10:     $o_j \leftarrow F_{A,O}[i]$ , obtain current state  $\tilde{s}$   
11:     $(s_1, s_2), r_e \leftarrow ExecuteOption(\tilde{s}, o_j)$   
12:    append  $r_e$  into  $R[(s_1, s_2)]$ ,  $q \leftarrow q + r_e$   
13:   end for  
14:   while env isn't terminal do  
15:     obtain current state  $\tilde{s}$   
16:      $(s_1, s_2), r_e \leftarrow ExecuteOption(\tilde{s}, o_G)$   
17:     if  $(s_1, s_2)$  not in  $R$  then  
18:        $R[(s_1, s_2)] \leftarrow list(r_e)$   
19:     else  
20:       append  $r_e$  into  $R[(s_1, s_2)]$   
21:     end if  
22:   end while  
23:   train options in  $O$  and calculate  $sr$   
24: end for
```

options (o_0, o_1, \dots, o_n) . If all action models in Π successfully finish, which indicates the chosen symbolic options are executed sequentially and termination conditions are satisfied, then the meta-controller would choose the global option o_G to explore the environment thoroughly. For each option $o_i = (pre_i, \pi_i, \beta_i)$, we refer to (Lyu et al. 2019) to design intrinsic rewards:

$$r_i(\tilde{s}) = \begin{cases} \phi & \beta_i(\tilde{s}) = True \\ r & otherwise \end{cases} \quad (5)$$

where ϕ is a constant and r is the reward gained from the environments when reach state \tilde{s} .

Planning and Learning

As shown in Algorithm 1, we firstly initialize an option set O only including o_G , an empty action model set A , an empty dictionary mapping symbolic state pairs to their external rewards R , an empty action models success ratio set sr and an empty plan Π_0 . When an episode t begins, we first get a start state \tilde{s}_0 from environment. Then we compute the symbolic initial state I by F and record the best plan Π^* , which is the plan generate in the last episode. Then meta-controller updates action models A , symbolic options set O , their mapping function $F_{A,O}$ and the planning goal G . Given current action models A and planning goal G , Metric-FF

planner (Hoffmann 2002) generates a new plan Π_t whose quality is higher than the last plan Π_{t-1} . If Π_t is empty, which indicates Metric-FF couldn't find a solution to solve the problem, we let $\Pi_t = \Pi^*$.

As for each action model a_i in plan Π_t , meta-controller chooses a corresponding symbolic option o_j by $F_{A,O}$. Then the controller interacts with environment by performing Deep Q-Learning, executes the action chosen by o_j 's inner policy and stores experience into o_j 's replay buffer until o_j terminates. After that, we get o_j 's initial symbolic state s_1 and a terminate symbolic state s_2 and an extrinsic reward r_e . In this way, we compute symbolic state pairs and their extrinsic rewards one by one and record these mappings by a dictionary R . Finally, quality q of plan Π_t is defined as the accumulated sum of extrinsic rewards.

If the environment isn't finished after executing Π_t , the meta-controller chooses the global option o_G to explore new symbolic states pairs in the environment. o_G stops exploring when the computed symbolic state changes and we calculate a symbolic state pair (s_1, s_2) and its external reward r_e . If (s_1, s_2) is a new symbolic state pair, we add it into R . This process repeats until the environment is terminated. Finally, when an episode ends, we train options in O and calculate success ratio for each action model.

Experiment

In this section, we evaluate our approach on two domains, Office World and Montezuma's Revenge in terms of data-efficiency, interpretability and transferability.


Office World

We first evaluate our approach on the Office World (Icarte et al. 2018) which is a simple multitask environment. In this environment, being initialized at a random location, the agent can move towards one of the four cardinal directions. Actions are valid only if the movement does not go through a wall. The agent can pick up cups of coffee or mails when it reaches the cell marked with blue cups or green envelopes, respectively. He can deliver coffee or mail to the office by reaching the cell marked with a purple hand. The symbol * means the place where the agent can not stay or reach.


Setup In this environment, the start location of the agent is randomly initialized at every episode. The agent is required to finish three tasks. The first and the second are to deliver a cup of coffee or a piece of mail to the office while the third is to hand both objects to the office. We compared SORL to h-DQN, a goal based HRL approach (Kulkarni et al. 2016). Since the state and action space are finite, we choose to implement these two approaches with q-table in both high and low levels.

Results We evaluate our approaches in terms of data-efficiency, interpretability and transferability.

- **Data-efficiency** In order to validate the data-efficiency, we train these two approaches in the three tasks and compare the corresponding performance at the same interaction steps. To demonstrate the transferability, we train the agent in task 3 along with the options learned in tasks 1



(a) Office World




(b) Learning Curve

task	proposition
task1	(haveCoffee), (deliveredCoffee)
task2	(haveEmail), (deliveredEmail)
task3	(haveCoffee), (deliveredCoffee) (haveEmail), (deliveredEmail)

(c) Propositions

task	action	description	option	effect
1	act0	get coffee	opt0	(haveCoffee)
	act1	deliver coffee	opt1	(not (haveCoffee)) (deliveredCoffee)
2	act0	get mail	opt2	(haveMail)
	act1	deliver mail	opt3	(not (haveMail)) (deliveredMail)
3	act0	get coffee	opt0	(haveCoffee)
	act1	get mail	opt1	(haveMail)
	act2	deliver coffee and mail	opt4	(not (haveCoffee)) (not (haveMail)) (deliveredCoffee) (deliveredMail)

(d) Learned Action Models and Options



(e) Learning Curve

Figure 4: Experimental Results in the Office World

and 2. To implement our approach, we design the propositions as shown in Fig. 4(c). As shown in Fig.4(b), from Task1 to Task3, SORL can get rewards faster than HRL.

- **Interpretability** Fig.4(d) shows action models and symbolic options learned in each task. Those action models describe the reason of making decisions at each step in a human understandable way. For example, we can explicitly know *act1* in task1 can be executed when the agent gets coffee and does not deliver it to the office, and the agent would deliver the coffee to the office and get a reward of 100 when *act1* is executed.
- **Transferability** By utilizing the options learned in tasks 1 and 2, we test the transferability of SORL and H-DQN in Task 3 and denote them as $SORL_t$ and HRL_t . As shown in Fig4(c), the performance of SORL and HRL is improved when transferring the learned knowledge. It verifies that compared to SORL, the converging speed of $SORL_t$ improves dramatically with only half of samples. We conjecture that the SORL is able to transfer the learned knowledge into other unseen environments.


Montezuma’s Revenge

Montezuma’s Revenge is an Atari game with sparse and delayed rewards. It requires the player to navigate through several rooms while collecting treasures. We conduct our experiments based on the first room shown in Fig.5(a). In this room, the player only obtains positive rewards when it fetches the key (+100) or opens a door (+300). Otherwise, the player would not receive any reward signal. The optimal solution is to climb down the ladders to obtain the key, then

return back to the platform and open a door, resulting in a maximum reward (+400).

Setup We compare our approach with HRL (Kulkarni et al. 2016) and SDRL (Lyu et al. 2019) as baselines, where SDRL is an approach that combines symbolic planning and RL with excellent results in complex environments with sparse rewards. SORL can automatically learn the action models while they are pre-defined by experts in SDRL. Besides, the option model can correspond to multiple action models in SORL while one in SDRL. We implement these approaches under an option-based HRL framework. In terms of the low level, we follow the network architecture used in (Kulkarni et al. 2016) and train this network with double-Q learning (van Hasselt, Guez, and Silver 2016) and prioritized experience replay (Schaul et al. 2016). Besides, both SORL and SDRL use a planner to generate high level policy while HRL utilizes a neural network. The intrinsic reward follows 5 with $\phi = 100$. The maximum steps in an episode and the threshold of success rate are set to be 500 and 0.95, respectively. To describe the environment, we abstract four local propositions (e.g., MiddleLadder, RightDoor, LeftLadder and RightLadder) and an object (Key).


Results We present the experimental results in Fig.5. It is evident that SORL can achieve the maximum reward (+400) in 0.7M samples while both SDRL and HRL need more than 1.5M samples, indicating the superior data-efficiency of SORL. However, to pick up the key (reward +100), SORL needs to interact with the environment with more than 0.3M steps, at which SDRL and HRL fall into the local optimum. This is because SORL randomly explores symbolic options



(a) Learning Curve

act	action model	
act1	pre	(atMiddleLadder)(keyexist)
	eff	(not (atMiddleLadder))(atRightLadder) (increase (quality) 0)
act5	pre	(atLeftLadder) (keyexist)
	eff	(not (keyexist)) (increase (quality) 100)
act10	pre	(atMiddleLadder) (not (keyexist))
	eff	(not (atMiddleLadder)) (atRightDoor) (increase (quality) 300)

(b) Examples of Learned Action Models



(c) learned Symbolic Options

action	description	option
act0	move from MiddleLadder to RightLadder without Key	opt0
act1	move from MiddleLadder to RightDoor without Key	opt1
act2	move from RightDoor to MiddleLadder without Key	opt2
act3	move from RightLadder to MiddleLadder without Key	opt3
act4	move from RightLadder to LeftLadder without Key	opt4
act5	pick up the Key	opt5
act6	move from LeftLadder to RightLadder without Key	opt6
act7	move from LeftLadder to RightLadder with Key	opt6
act8	move from RightLadder to MiddleLadder with Key	opt3
act9	move from RightLadder to LeftLadder with Key	opt4
act10	move from MiddleLadder to RightDoor with Key	opt1

(d) Learned Action Models

Figure 5: Experimental Results in Montezuma’s Revenge

and it is easier to find options closer to the starting point. After finding these options, SORL would train them sequentially instead of directly learning options on the path of getting the key. One option model corresponds to one action model in SDRL while several action models in SORL. The ability of reusing the learned symbolic options enables SORL to converge faster than SDRL. Take Fig.5(c) as an example, the *opt1* representing the move from middle ladder to right door, is firstly trained at the beginning when the player does not get the key. After the player picks up the key, SORL only needs a small amount of data to fine-tune *opt1* when the player moves from the middle ladder to right door with a key. However, both SDRL and HRL start training the options after the player moves to the middle ladder with a key, consuming more interaction resources. Different from option-based HRL and SDRL, SORL can learn the initial and termination condition of symbolic options automatically. The action models used in SDRL need to be constructed by human in advance while they are learned from the trajectories in SORL, saving labour resources. We present some of the learned action models in Fig.5(b) and the effects of symbolic options in Fig.5(c). Fig. 5(b) describes the preconditions and effects of each action model and we can see that if the player is at LeftLadder and the key exists, then the player can obtain a key and reward (+100) by executing action5. Fig. 5(c) shows the learned options in SORL and the order of options actually

does not match the optimal order because SORL randomly explore the environment and options 0-3 are easier to learn. We describe the meaning of all learned action models and their corresponding options in Fig. 5(d). It is easy to see that act7 to act10 correspond to the options explored before, so these options can be reused to improve the data-efficiency.

Conclusions

In this paper, we propose a novel framework SORL to study a creativity property of an AI agent, which can automatically “accumulate” new action models and symbolic options from trajectories, and leverage the accumulated action models and symbolic options to instruct RL to explore efficiently in environments with sparse and delayed rewards. Compared with previous approaches, our experimental results exhibit that SORL has better sampling efficiency. Besides, SORL requires less prior knowledge and is able to give better interpretability and transferability with the learnt action models and symbolic options.

In the future, it would be interesting to investigate the possibility of learning more expressive planning models, such as learning Hierarchical Task Networks (Zhuo and Yang 2014) and PDDL models (Zhuo et al. 2010), as well as integrating different learning mechanisms, such as transfer learning (Zhuo and Yang 2014; Shen et al. 2020).

Acknowledgements

This research was funded by the National Natural Science Foundation of China (Grant No. 62076263), Guangdong Natural Science Funds for Distinguished Young Scholar (Grant No. 2017A030306028), Guangdong special branch plans young talent with scientific and technological innovation (Grant No. 2017TQ04X866), Pearl River Science and Technology New Star of Guangzhou and Guangdong Province Key Laboratory of Big Data Analysis and Processing.

References

- Aradi, S. 2020. Survey of Deep Reinforcement Learning for Motion Planning of Autonomous Vehicles. *CoRR*, abs/2001.11231.
- Boden, M. A. 1998. Creativity and artificial intelligence. *Artificial Intelligence*, 103(1): 347–356.
- Boden, M. A. 2009. Computer Models of Creativity. *AI Mag.*, 30(3): 23–34.
- Hoffmann, J. 2002. Extending FF to Numerical State Variables. In van Harmelen, F., ed., *Proceedings of the 15th European Conference on Artificial Intelligence, ECAI 2002, Lyon, France, July 2002*, 571–575. IOS Press.
- Icarte, R. T.; Klassen, T. Q.; Valenzano, R. A.; and McIlraith, S. A. 2018. Advice-Based Exploration in Model-Based Reinforcement Learning. In Bagheri, E.; and Cheung, J. C. K., eds., *Advances in Artificial Intelligence - 31st Canadian Conference on Artificial Intelligence, Canadian AI 2018, Toronto, ON, Canada, May 8-11, 2018, Proceedings*, volume 10832 of *Lecture Notes in Computer Science*, 72–83. Springer.
- Illanes, L.; Yan, X.; Icarte, R. T.; and McIlraith, S. A. 2020. Symbolic Plans as High-Level Instructions for Reinforcement Learning. In Beck, J. C.; Buffet, O.; Hoffmann, J.; Karpas, E.; and Sohrabi, S., eds., *Proceedings of ICAPS*, 540–550. AAAI Press.
- James, S.; Rosman, B.; and Konidaris, G. 2020. Learning Portable Representations for High-Level Planning. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, 4682–4691. PMLR.
- Kulkarni, T. D.; Narasimhan, K.; Saeedi, A.; and Tenenbaum, J. 2016. Hierarchical Deep Reinforcement Learning: Integrating Temporal Abstraction and Intrinsic Motivation. In Lee, D. D.; Sugiyama, M.; von Luxburg, U.; Guyon, I.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, 3675–3683.
- Lee, J.; Katz, M.; Agravante, D. J.; Liu, M.; Klinger, T.; Campbell, M.; Sohrabi, S.; and Tesauro, G. 2021. AI Planning Annotation in Reinforcement Learning: Options and Beyond. In *Planning and Reinforcement Learning PRL Workshop at ICAPS*.
- Leonetti, M.; Iocchi, L.; and Stone, P. 2016. A synthesis of automated planning and reinforcement learning for efficient, robust decision-making. *Artif. Intell.*, 241: 103–130.
- Lyu, D.; Yang, F.; Liu, B.; and Gustafson, S. 2019. SDRL: Interpretable and Data-Efficient Deep Reinforcement Learning Leveraging Symbolic Planning. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, 2970–2977. AAAI Press.
- Martínez, D.; Alenyà, G.; Torras, C.; Ribeiro, T.; and Inoue, K. 2016. Learning Relational Dynamics of Stochastic Domains for Planning. In Coles, A. J.; Coles, A.; Edelkamp, S.; Magazzeni, D.; and Sanner, S., eds., *Proceedings of the Twenty-Sixth International Conference on Automated Planning and Scheduling, ICAPS 2016, London, UK, June 12-17, 2016*, 235–243. AAAI Press.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *nature*, 518(7540): 529–533.
- Ng, J. H. A.; and Petrick, R. P. A. 2019. Incremental Learning of Planning Actions in Model-Based Reinforcement Learning. In Kraus, S., ed., *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, 3195–3201. ijcai.org.
- Ryan, M. R. K. 2002. Using Abstract Models of Behaviours to Automatically Generate Reinforcement Learning Hierarchies. In Sammut, C.; and Hoffmann, A. G., eds., *Machine Learning, Proceedings of the Nineteenth International Conference (ICML 2002), University of New South Wales, Sydney, Australia, July 8-12, 2002*, 522–529. Morgan Kaufmann.
- Sarathy, V.; Kasenberg, D.; Goel, S.; Sinapov, J.; and Scheutz, M. 2020. SPOTTER: Extending Symbolic Planning Operators through Targeted Reinforcement Learning. *arXiv preprint arXiv:2012.13037*.
- Schaul, T.; Quan, J.; Antonoglou, I.; and Silver, D. 2016. Prioritized Experience Replay. In Bengio, Y.; and LeCun, Y., eds., *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Shen, J.; Zhuo, H. H.; Xu, J.; Zhong, B.; and Pan, S. J. 2020. Transfer Value Iteration Networks. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, 5676–5683. AAAI Press.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; van den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; Dieleman, S.; Grewe, D.; Nham, J.; Kalchbrenner, N.; Sutskever, I.; Lillicrap, T. P.; Leach, M.; Kavukcuoglu, K.; Graepel, T.; and Hassabis, D. 2016. Mastering the game of Go with deep neural networks and tree search. *Nat.*, 529(7587): 484–489.

- Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; Bolton, A.; et al. 2017. Mastering the game of Go without human knowledge. *Nature*, 550(7676): 354–359.
- Sutton, R. S.; Precup, D.; and Singh, S. P. 1999. Between MDPs and Semi-MDPs: A Framework for Temporal Abstraction in Reinforcement Learning. *Artif. Intell.*, 112(1-2): 181–211.
- van Hasselt, H.; Guez, A.; and Silver, D. 2016. Deep Reinforcement Learning with Double Q-Learning. In Schuurmans, D.; and Wellman, M. P., eds., *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, 2094–2100. AAAI Press.
- Yang, F.; Lyu, D.; Liu, B.; and Gustafson, S. 2018. PEORL: Integrating Symbolic Planning and Hierarchical Reinforcement Learning for Robust Decision-Making. In Lang, J., ed., *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, 4860–4866. ijcai.org.
- Yang, Q.; Wu, K.; and Jiang, Y. 2007. Learning action models from plan examples using weighted MAX-SAT. *Artif. Intell.*, 171(2-3): 107–143.
- Zhou, Z.; Li, X.; and Zare, R. N. 2017. Optimizing chemical reactions with deep reinforcement learning. *ACS central science*, 3(12): 1337–1344.
- Zhuo, H. H.; and Kambhampati, S. 2017. Model-lite planning: Case-based vs. model-based approaches. *Artif. Intell.*, 246: 1–21.
- Zhuo, H. H.; and Yang, Q. 2014. Action-model acquisition for planning via transfer learning. *Artif. Intell.*, 212: 80–103.
- Zhuo, H. H.; Yang, Q.; Hu, D. H.; and Li, L. 2010. Learning complex action models with quantifiers and logical implications. *Artif. Intell.*, 174(18): 1540–1569.