

Refining Manually-Designed Symbol Grounding and High-Level Planning by Policy Gradients

Takuya Hiraoka¹², Takashi Onishi¹², Yoshimasa Tsuruoka²³,

¹ NEC Central Research Laboratories

² National Institute of Advanced Industrial Science and Technology

³ The University of Tokyo

t-hiraoka@ce.jp.nec.com, t-onishi@bq.jp.nec.com, tsuruoka@logos.t.u-tokyo.ac.jp

Abstract

Hierarchical planners that produce interpretable and appropriate plans are desired, especially in its application to supporting human decision making. In the typical development of the hierarchical planners, higher-level planners and symbol grounding functions are manually created, and this manual creation requires much human effort. In this paper, we propose a framework that can automatically refine symbol grounding functions and a high-level planner to reduce human effort for designing these modules. In our framework, symbol grounding and high-level planning, which are based on manually-designed knowledge bases, are modeled with semi-Markov decision processes. A policy gradient method is then applied to refine the modules, in which two terms for updating the modules are considered. The first term, called a reinforcement term, contributes to updating the modules to improve the overall performance of a hierarchical planner to produce appropriate plans. The second term, called a penalty term, contributes to keeping refined modules consistent with the manually-designed original modules. Namely, it keeps the planner, which uses the refined modules, producing interpretable plans. We perform preliminary experiments to solve the Mountain car problem, and its results show that a manually-designed high-level planner and symbol grounding function were successfully refined by our framework.

1 Introduction

Hierarchical planners have been widely researched in artificial intelligence communities. One of the main reasons for that is that the hierarchical planners can divide complex planning problems, which flat planners cannot solve, into a series of more simple sub-problems, by using high-level knowledge about the planning problem (e.g., [Nilsson, 1984; Choi and Amir, 2009; Kaelbling and Lozano-Pérez, 2011]).

A hierarchical planner is composed of multiple planner layers that are typically divided into two types: high-level and low-level. A low-level planner performs micro-level planning, and it deals with raw information about an environment.

In contrast, a high-level planner performs macro-level planning, and it deals with more abstract symbolic information. The raw and abstract symbolic information are mapped to each other by *symbol grounding* functions. Imagine that a hierarchical planner is used for controlling a humanoid robot to put a lemon on a board. Here, the high-level planner makes a plan such as “Pick a lemon up, and then put it on a board.” The low-level planner makes a plan for controlling the robot’s motors according to sensor inputs, to achieve sub-goals given by the high-level planner (e.g., “Pick a lemon up”). As the low-level planner cannot understand what “Pick a lemon up” means, the symbol ground function converts it into actual values, in the environment, which the low-level planner can understand.

Hierarchical planners are often used for supporting human decision making (e.g., in supply chain [Özdamar *et al.*, 1998] or clinical operations [Fdez-Olivares *et al.*, 2011]). In such cases, people make decisions on the basis of a plan, and thus it is necessary that 1) they understand the plan (especially one of a high-level planner) and 2) they can reach satisfying outcomes by following the plan (i.e., the hierarchical planner gives appropriate plans).

In many previous studies on hierarchical planners, symbol grounding functions and high-level planners were designed manually [Nilsson, 1984; Malcolm and Smithers, 1990; Cambon *et al.*, 2009; Choi and Amir, 2009; Dornhege *et al.*, 2009; Wolfe *et al.*, 2010; Kaelbling and Lozano-Pérez, 2011]. Although this makes it possible for people to understand the plans easily, much human effort is needed to carefully design a hierarchical planner that provides appropriate plans.

Konidaris *et al.* [2014; 2015; 2016] have proposed frameworks for automatically constructing symbol grounding functions and high-level planners, but they require a human to carefully analyze them to understand the plans. These constructed modules are often complicated and, in such cases, analysis becomes a burden.

In this paper, we propose a framework that automatically refines manually-designed symbol grounding functions and high-level planners, with a policy gradient method. Our framework differs from frameworks proposed in the aforementioned previous studies on the basis of the following points:

- Unlike the hierarchical planners based solely on manually-designed symbol grounding functions and

high-level planners [Nilsson, 1984; Malcolm and Smithers, 1990; Cambon *et al.*, 2009; Choi and Amir, 2009; Dornhege *et al.*, 2009; Wolfe *et al.*, 2010; Kaelbling and Lozano-Pérez, 2011], our framework refines these modules without human intervention. This automated refinement reduces the design workload for the modules.

- Unlike the frameworks that automatically construct symbol grounding functions and high-level planners [Konidaris *et al.*, 2014; 2015; Konidaris, 2016], our framework refines these while attempting to keep the resulting symbol grounding consistent with prior knowledge of the definition of the symbols as much as possible (see Section 4). Therefore, a person can understand the plan that high-level planners output without careful analysis of the refined modules.

In this paper, we first explain our hierarchical planner (including the high-level planner and symbol grounding functions), and how these are designed (Section 3). Then, we introduce the framework designed to refine them (Section 4). Finally, we experimentally demonstrate the effectiveness of our framework (Section 5).

2 Preliminaries

Our framework, introduced in Section 4, is based on semi-Markov decision processes (SMDPs) and policy gradient methods.

2.1 Semi-Markov Decision Processes

SMDPs are a framework for modeling a decision problem in an environment where a sojourn time in each state is a random variable, and it is defined as a tuple $\langle S, O, R, P, \gamma \rangle$. $S \subseteq \mathbb{R}^n$ is the n -dimensional continuous state space; $O(s)$ is a function that returns a finite set of options [Sutton *et al.*, 1999] available in the environment’s state $s \in S$; $R(s', t|s, o)$ is the reward received when option $o \in O(s)$ is executed at s ; arriving in state $s' \in S$ after t time steps; $P(s', t|s, o)$ is a probability of $s' \in S$, and t after executing o in s ; and $\gamma \in [0, 1]$ is a discount factor.

Given SMDPs, our interest is to find an optimal policy over options $\pi^*(o|s)$:

$$\pi^* = \arg \max_{\pi} V_{\pi}(s_0), \quad (1)$$

$$V_{\pi}(s_0) = E_{\pi}[R(s_1, t_0|s_0, o_0) + \gamma R(s_2, t_1|s_1, o_1) + \dots] \quad (2)$$

where (s_0, o_0, t_0, s_1) and (s_1, o_1, t_1, s_2) are transitions of a state, an option, time steps elapsed while executing the option, and the arriving state after executing the option.

2.2 Policy Gradient

To find π^* , we use a policy gradient method [Sutton *et al.*, 2000]. In a policy gradient method, a policy $\pi_{\theta}(o|s)$ parameterized by θ is introduced to approximate π^* , and the approximation is performed by updating θ with a gradient. Although there are many policy gradient implementations (e.g., [Kakade, 2002; Silver *et al.*, 2014; Schulman *et al.*, 2015]),

we use REINFORCE [Williams, 1992]. In REINFORCE, θ is updated as follows:

$$\theta \leftarrow \theta + \alpha \nabla_{\theta} \log \pi_{\theta}(\tilde{s}_0, \tilde{o}_0) V_{\pi_{\theta}}(\tilde{s}_0), \quad (3)$$

$$V_{\pi_{\theta}}(\tilde{s}_0) = R(\tilde{s}_1, \tilde{t}_0|\tilde{s}_0, \tilde{o}_0) + \gamma \tilde{R}(\tilde{s}_2, \tilde{t}_1|\tilde{s}_1, \tilde{o}_1) + \dots, \quad (4)$$

where α is a learning rate and $(\tilde{s}_0, \tilde{o}_0, \tilde{t}_0, \tilde{s}_1)$ and $(\tilde{s}_1, \tilde{o}_1, \tilde{t}_1, \tilde{s}_2)$ are transitions of state, the executing option, elapsed time steps, and arriving state, which are sampled on the basis of π_{θ} in a time horizon. Other variables and functions are the same as those introduced in Section 2.1. We decided to use REINFORCE for our work because it has successfully worked in recent work [Silver *et al.*, 2016; Das *et al.*, 2017].

3 Hierarchical Planner with Symbol Grounding Functions

In this section, we first describe the outline of a hierarchical planner (including the high-level planner) with symbol grounding functions, which are manually designed. We then provide concrete examples of them. The high-level planner and symbol grounding functions described here are refined by the framework, which is proposed in Section 4.

The hierarchical planner (Figure 1) is composed of two symbol grounding functions (one for abstraction and the other for concretization), a high-level planner, a low-level planner, and two knowledge bases (one each for the high-level and low-level planners). These modules work as follows:

Step 1 : The symbol grounding function for abstraction receives raw information, abstracts it to a symbolic information on the basis of its knowledge base, and then outputs the symbolic information.

Step 2 : The high-level planner receives the abstract symbolic information, makes a plan using its knowledge base, and then outputs abstract symbolic information as a sub-goal, which indicates the next abstract state to be achieved.

Step 3 : The symbol grounding function for concretization receives the abstract symbolic information, concretizes it to raw information about a sub-goal, which specifies an actual state to be achieved, then outputs the raw information on the sub-goal. This module performs the concretization on the basis of its the knowledge base.

Step 4 : The low-level planner receives the raw information about a sub-goal and then interacts with the environment to achieve the given sub-goal. In the interaction, the low-level planner outputs primitive actions in accordance with the raw information given by the environment. The interaction continues until the low-level planner achieves the given sub-goal, or until the total number of elapsed time steps reaches a given threshold.

Step 5 : If the raw information from the environment is not a goal or terminal state, return to the Step 1.

The knowledge bases for symbol grounding functions and the high-level planners are designed manually.

Knowledge base for high-level planners is described as a simple planning domain definition language (PDDL) [McDermott *et al.*, 1998]. In a PDDL, objects, predicates, goals, and operators are manually specified.

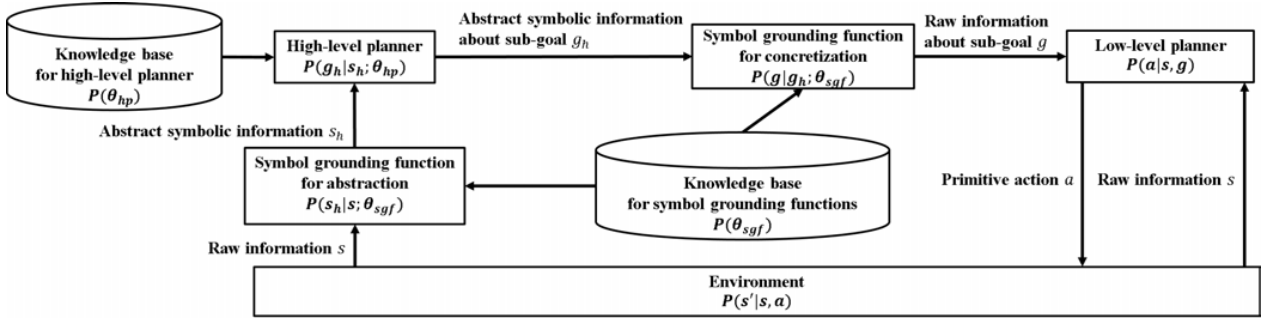


Figure 1: Outline of hierarchical planner with grounding functions.




Figure 2: Mountain car with abstract symbols.

The objects and predicates are for building logical formulae, which specify the possible states in the planning domain. The operators are represented as a pair of preconditions and effects. The preconditions represent the states required for applying the operator, and the effects represent the arriving states after applying the operators. We use PDDLs in this work because they are widely used for describing a knowledge base for symbolic planners.

Knowledge base for symbol grounding functions is described as a list of maps between abstract symbolic information and corresponding raw information. In this paper, to simplify the problem, we assume that each item of abstract symbolic information is mapped into one interval of raw information. Despite its simplicity, it is useful for representing, for example, typical spatial information.

Here, we describe the knowledge bases and how the hierarchical planner works to solve the mountain car problem [Moore, 1991] (Figure 2). In this problem, a car is placed within a deep valley, and its goal is to drive out by going up the right side hill. However, as the car’s engine is not strong enough, it needs to first drive back and forth between the two hills to generate momentum. In this problem, the hierarchical planner receives raw information (the position and velocity of the car) from the environment and is required to make a plan to move it to the goal (the top of the right side hill).

An example of knowledge for the high-level planner is shown in Table 1. In this example, objects are composed of only a “Car.” Predicates are composed of four instances (“Bottom_of_hills(x), On_right_side_hill(x), On_left_side_hill(x), and At_top_of_right_side_hill(x)”). For example, “On_right_side_hill(Car)” means that the car is on

Table 1: Example knowledge for high-level planners. Upper part describes examples of objects, predicates, and goals. Lower part describes examples of operators.

| | |
|------------|--|
| Objects | $x = \{\text{Car}\}$ |
| Predicates | Bottom_of_hills(x), On_right_side_hill(x), On_left_side_hill(x), At_top_of_right_side_hill(x) |
| Goals | At_top_of_right_side_hill(Car) |

| Operators | Preconditions | Effects |
|-----------|---------------------------|----------------------------------|
| Opr.1 | Bottom_of_hills(x) | On_right_side_hill(x) |
| Opr.2 | On_right_side_hill(x) | On_left_side_hill(x) |
| Opr.3 | On_left_side_hill(x) | At_top_of_right_side_hill(x) |

Table 2: Example knowledge for symbol grounding functions.

| Abstract symbolic informations | Interval of raw information |
|--------------------------------|-----------------------------|
| Bottom_of_hills(Car) | position $\in [-0.6, -0.4]$ |
| On_right_side_hill(Car) | position $\in [-0.2, 0.4]$ |
| On_left_side_hill(Car) | position $\in [-1.2, -0.8]$ |
| At_top_of_right_side_hill(Car) | position $\in [0.6, 0.8]$ |

the right side hill. Operators are composed of three types that refer to a transition of the objects on the hills. For example, “Opr.1” refers to the transition that object x has moved from the bottom of the hills to the right side hill.

An example of the knowledge for symbol grounding functions is shown in Table 2. This example shows mappings between abstract symbolic information (the location of the car), and corresponding intervals of raw information (the actual value of the car’s position). For example, “Bottom_of_hills(Car)” is mapped to the position of the car is in the interval $[-0.6, -0.4]$.

Given the knowledge described in Tables 1 and 2, an example of how the hierarchical planner works is shown as follows:

Example of Step 1: The symbol grounding function for abstraction receives raw information (position=-0.5 and velocity=0). The position is in the interval $[-0.6, -0.4]$, which corresponds to the “Bottom_of_hills(Car)” in Table 2. Therefore, the symbol grounding function outputs “Bottom_of_hills(Car).”

Example of Step 2: The high-level planner receives “Bottom_of_hills(Car),” and makes a plan to achieve the goal (“At_top_of_right_side_hill(Car)”). By using the knowledge in Table 1, the high-level planner makes the plan [Bottom_of_hills(Car) \rightarrow

On_right_side_hill(Car) \rightarrow On_left_side_hill(Car) \rightarrow At_top_of_right_side_hill(Car)], which means “Starting at the bottom of the hills, visit, in order, the right side hill, the left side hill, and the top of the right side hill.” After following the plan, the high-level planner outputs “On_right_side_hill(Car).”

Example of Step 3: The symbol grounding function receives “On_right_side_hill(Car),” and concretizes it to raw information about sub-goal (position= 0.1, velocity=*). Here, the position in the raw information is determined as the mean of the corresponding interval $[-0.2, 0.4]$ in Table 2. In addition, the mask (represented by “*”) is putted to filter out factors in raw information, which is irrelevant in the sub-goal (i.e., velocity in this example).

Example of Step 4: The low-level planner receives position= 0.1 and the mask. To move the car to the given sub-goal (position=0.1), the low-level planner makes a plan to accelerate the car. This planning is performed by model predictive control [Camacho and Alba, 2013]. The low-level planner terminates itself when the car arrives at the given sub-goal (position=0.1), or when it takes a primitive action 20 times.

4 Framework for Refining Grounding Function and High Level Planner

In this section, we propose a framework for refining the symbol grounding functions and the high-level planner introduced in the previous section. In our framework, symbol grounding and high-level planning, which are based on manually-designed knowledge bases, are modeled with SMDPs. Refinement of the symbol grounding functions and the high-level planner is achieved by applying policy gradients to the model. First, we introduce an abstract model and then provide an example of its implementation in the mountain car problem. Finally, we explain how the policy gradient method is applied to the model.

4.1 Modeling Symbol Grounding and High-Level Planning with SMDPs

We model symbol grounding and high-level planning, which are based on manually-designed knowledge bases, with SMDPs. The symbol grounding functions and the high-level planner are modeled as components of the parameterized policy. In addition, the knowledge bases are modeled as priors for the policy’s parameters.

We first assume that information and modules, which appear in hierarchical planning, are represented as random variables and probability functions, respectively (Figure 1). Suppose that S_h is a set of all possible symbols the symbol grounding functions and the high-level planner deal with, raw information is represented as an n -dimensional vector, and $A \subset \mathbb{R}^m$ is a set of all possible primitive actions. We denote raw information by $s, s' \in \mathbb{R}^n$ ¹, abstract symbol information by $s_h \in S_h$, abstract symbolic information about a sub-goal

¹The denotation is the same as that of the state described in Section 2.1 because the raw information is modeled as the state.




Figure 3: SMDPs for our framework.

by $g_h \in S_h$, raw information about a sub-goal $g \in \mathbb{R}^n$, and a primitive action by $a \in A$. In addition, we denote the symbol grounding function for abstraction by $P(s_h|s; \theta_{sgf})$, the symbol grounding function for concretization by $P(g|g_h; \theta_{sgf})$, the high-level planner by $P(g_h|s_h; \theta_{hp})$, the low-level planner by $P(a|s, g)$, the environment by $P(s'|s, a)$, the knowledge base for high-level planners by $P(\theta_{hp})$, and the knowledge base for the high-level planner by $P(\theta_{sgf})$. Here, θ_{sgf} and θ_{hp} are the parameters for the symbol grounding functions and the high-level planner, respectively.

High-level planning and symbol grounding based on the knowledge base are modeled as SMDPs (Figure 3). In this model, the components of SMDPs (i.e., an option, a state, a reward, and a transition probability) are implemented as follows:

Option o : is implemented as a tuple $\langle s_h, g_h, g \rangle$ of abstract symbolic information s_h , abstract symbolic information about a sub-goal g_h , and raw information about a sub-goal g .

State s : is implemented as raw information.

Reward $R(s', t|s, o)$: is the cumulative reward given by the environment $P(s'|s, a)$, while the low-level planner $P(a|s, g)$ is interacting with $P(s'|s, a)$.

Transition probability $P(s', t|s, o)$: is implemented as a function, which represents the state transition proceeded by the interaction between the low-level planner $P(a|s, g)$ and the environment $P(s'|s, a)$. Note that although the transition probability receives option $\langle s_h, g_h, g \rangle$, only g is used in the transition probability.

In this model, the parameterized policy π_θ is implemented to control abstraction of raw information, high-level planning, and concretization of abstract symbolic information, in accordance with the knowledge bases. Formally, π_θ is implemented as follows:

$$\begin{aligned} \pi_\theta(g, g_h, s_h|s) &= P(g, g_h, s_h, \theta_{sgf}, \theta_{hp}|s) \\ &= P(g|g_h; \theta_{sgf})P(g_h|s_h; \theta_{hp})P(s_h|s; \theta_{sgf}) \\ &\quad \cdot P(\theta_{sgf})P(\theta_{hp}). \end{aligned} \quad (5)$$

The right term in the second line can be derived by decomposing the joint probability in the first line, in accordance with the probabilistic dependency shown in Figure 2. Note that, in this equation, θ is represented as $\theta_{sgf}||\theta_{hp}$, i.e., a concatenation of θ_{sgf} and θ_{hp} . By using this representation for π_θ , we can derive an update expression, which can refine θ_{sgf} and θ_{hp} keeping them consistent with $P(\theta_{sgf})$ and $P(\theta_{hp})$. See Section 4.3 for details.

$P(\theta_{sgf})$ and $P(\theta_{hp})$ are needed to reflect the manually-designed knowledge bases. To do so, first, $P(\theta_{sgf})$ and $P(\theta_{hp})$ are implemented as parametric distributions $P(\theta_{sgf}; \theta'_{sgf})$ and $P(\theta_{hp}; \theta'_{hp})$, respectively, and their hyper-parameters θ'_{sgf} and θ'_{hp} are determined to replicate manually-designed symbol grounding functions and high-level planners. More formally, we use θ'^{*}_{sgf} and θ'^{*}_{hp} as the optimal parameters of θ'_{sgf} and θ'_{hp} , respectively, acquired by the following equations:

$$\theta'^{*}_{sgf} = \arg \min_{\theta'_{sgf}} D_{sgf} \left(P(g|g_h; \theta_{sgf}) P(\theta_{sgf}; \theta'_{sgf}) \right), \quad (6)$$

$$\theta'^{*}_{hp} = \arg \min_{\theta'_{hp}} D_{hp} \left(P(g_h|s_h; \theta_{hp}) P(\theta_{hp}; \theta'_{hp}) \right), \quad (7)$$

where D_{sgf} and D_{hp} are a divergence (e.g., KL divergence) from the manually-designed symbol grounding function and high-level planner, respectively. D_{sgf} and D_{hp} are abstract criteria, and thus, there are many implementations of functionals “ $\arg \min_{\theta'_{sgf}} D_{sgf}(\cdot)$ ” and “ $\arg \min_{\theta'_{hp}} D_{hp}(\cdot)$ ”.

4.2 An Example of Model Implementation to Solve the Mountain Car Problem

We introduced an abstract model for symbol grounding and high-level planning with knowledge bases in the previous section. In this section, we provide an example of an implementation of the model to solve the mountain car problem.

First, S_h and A are implemented as follows:

$$S_h := \left\{ \begin{array}{l} \text{Bottom_of_hills}(\text{Car}), \text{On_right_side_hill}(\text{Car}), \\ \text{On_left_side_hill}(\text{Car}), \text{At_top_of_right_side_hill}(\text{Car}) \end{array} \right\} \quad (8)$$

$$A := [-1.0, 1.0]. \quad (9)$$

S_h is implemented in accordance with the knowledge shown in Table 2. A is implemented in accordance with the definition of actions to solve the mountain car problem, and represented as a set of values for the acceleration of the car.

Second, the probabilities of the modules in the hierarchical planner are implemented as follows:

$$P(s_h|s; \theta_{sgf}) := \frac{N(s|\mu_{s_h}, \sigma_{s_h})}{\sum_{s'_h \in S_h} N(s|\mu_{s'_h}, \sigma_{s'_h})}, \quad (10)$$

$$P(g|g_h; \theta_{sgf}) := N(g|\mu_{g_h}, \sigma_{g_h}), \quad (11)$$

$$P(g_h|s_h; \theta_{hp}) := \frac{\exp(\phi(s_h) \mathbf{w}_{g_h}^T)}{\sum_{g'_h \in S_h} \exp(\phi(s_h) \mathbf{w}_{g'_h}^T)}. \quad (12)$$

$P(s_h|s; \theta_{sgf})$ is implemented as the normalized likelihood of a normal distribution (Eq. (10)), and $P(g|g_h; \theta_{sgf})$ is implemented as a normal distribution (Eq. (11)). In Eq. (10) and Eq. (11), $N(s|\mu_{s_h}, \sigma_{s_h})$ represents a normal distribution for s , which is parameterized by mean μ_{s_h} and standard deviation σ_{s_h} , s.t., $\forall s_h \in S_h$. Note that μ_{s_h} and σ_{s_h} are identical to μ_{g_h} and σ_{g_h} , respectively. $P(s_h|s; \theta_{sgf})$ is implemented as a softmax function (Eq. (12)). In Eq. (12), $\phi(s_h)$ is a base function that returns a one-hot vector $\in \{0, 1\}^{|S_h|}$ in which only one

element corresponding to the value of s_h is set to a value of 1, and the other elements are set to a value of 0. $\mathbf{w}_{g_h} \in \mathbb{R}^{|S_h|}$ is a weight vector, s.t., $\forall g_h \in S_h$. In this implementation, θ_{sgf} is a vector composed of μ_{s_h}, σ_{s_h} , s.t., $s_h \in S_h$, and θ_{hp} is a vector composed of the set \mathbf{w}_{g_h} , s.t., $g_h \in S_h$. $P(a|s, g)$ and $P(s'|s, a)$ are implemented as deterministic functions, which represent the simulator of environment² and the model predictive controller.

Third, the reward function $R(s', t|s, o)$ is implemented as follows:

$$R(s', t|s, o) := \sum_{i=0}^t \gamma^i r(s_i, a_i), \quad (13)$$

$$r(s_i, a_i) = \begin{cases} 100 & (\text{if car position in } s_i > 0.6) \\ -a_i & (\text{otherwise}) \end{cases} \quad (14)$$

where s_i and a_i are a state and a primitive action sampled from the environment i time steps later from the executing option o , respectively. Eq. (14) represents “low-level” reward $r(s_i)$, which is fed in accordance with a_i and the car position included in s_i .

Fourth, $P(\theta_{sgf}; \theta'_{sgf})$ and $P(\theta_{hp}; \theta'_{hp})$ are implemented as follows:

$$P(\theta_{sgf}; \theta'_{sgf}) := \prod_{s_h \in S_h} N(\mu_{s_h} | \mu'_{s_h}, 1) \cdot \text{Uni}(\sigma_{s_h}), \quad (15)$$

$$P(\theta_{hp}; \theta'_{hp}) := \prod_{g_h \in S_h} \prod_{i=0}^{|S_h|} N(w_{g_h, i} | \mu'_{w_{g_h, i}}, 1). \quad (16)$$

Eq. (15) represents a distribution for μ_{s_h} and σ_{s_h} . The component for μ_{s_h} is a normal distribution, which has mean μ'_{s_h} and standard deviation 1, and the component for σ_{s_h} is a uniform distribution $\text{Uni}(\sigma_{s_h})$. In addition, Eq. (16) represents the normal distribution for $w_{g_h, i}$, which is the i -th element of \mathbf{w}_{g_h} . This distribution has mean $\mu'_{w_{g_h, i}}$ and standard deviation 1. Note that, in this implementation, θ'_{sgf} and θ'_{hp} are μ'_{s_h} and $\mu'_{w_{g_h, i}}$, respectively.

Finally, functionals in Eq. (6) and Eq. (7), are implemented as follows:

Implementation of $\arg \min_{\theta'_{sgf}} D_{sgf}(\cdot)$: Using Eq. (15), μ'_{s_h} is

set as the mean of the corresponding interval, which is defined in the knowledge base for grounding functions. For example, $\mu'_{\text{Bottom_of_hills}(\text{car})}$ is determined as -0.5 , the mean of $[-0.6, -0.4]$ in Table 2.

Implementation of $\arg \min_{\theta'_{hp}} D_{hp}(\cdot)$: Using Eq. (7), $\mu'_{w_{g_h, i}}$ is

determined by Algorithm 1. The algorithm is outlined as follows: first initialize $\mu_{w_{g_h, i}}$ with val_{\min} (line 1–3), and if the operator, in which s_h refers to the preconditions and s'_h refers to the effects, is contained in knowledge base KB_{hp} , the corresponding weight is initialized with Val_{\min} (line 4–11). KB_{hp} is initialized in accordance with Table 1 before it is passed to the algorithm.

²Open AI gym was used as the simulator: <https://github.com/openai/gym/wiki/MountainCarContinuous-v0>

Algorithm 1 Implementation of $\arg \min D_{hp}$
 θ'_{hp}

Require: The following variables are given:

- (1) Set of abstract symbolic information S_h .
- (2) Set of operators K_{hp} included in the knowledge base for the high-level planner. Each operator is represented as a tuple (precondition, effect).
- (3) Set of hyper parameters $\mu'_{w_{gh}}$ for all possible abstract symbolic information s_h .
- (4) Weight value val_{in} to be assigned to the weight of an operator, which is included in the knowledge base.
- (5) Weight value val_{nin} to be assigned to the weight of an operator, which is not included in the knowledge base.
- (6) Index function $I(s_h)$ that maps s_h to index $i \in \mathbb{I}$. i used to access the element of $\mu'_{w_{gh}}$.

```

1: for  $s_h \in S_h$  do
2:   Initialize  $\mu'_{w_{gh}}$  with  $val_{nin}$ 
3: end for
4: for  $s_h \in S_h$  do
5:    $i \leftarrow 0$ 
6:   for  $s'_h \in S_h$  do
7:     if  $(s_h, s'_h) \in K_{hp}$  then
8:        $\mu'_{w_{gh}, I(s'_h)} \leftarrow val_{in}$ 
9:     end if
10:   end for
11: end for

```

4.3 Refining Symbol Grounding and High-Level Planning with Policy Gradients

Refining the high-level planner $P(g_h|s_h; \theta_{hp})$ and symbol grounding functions ($P(s_h|s; \theta_{sgf})$ and $P(g|g_h; \theta_{sgf})$) is achieved by a parameter update in Eq. (17). This equation contains two unique terms: a **reinforcement term** and a **penalty term**. The reinforcement term contributes to updating the parameters to maximize the expected cumulative reward, as in standard reinforcement learning. The penalty term contributes to keeping the parameters consistent with the priors (i.e., manually-designed knowledge bases). This update is derived by substituting $\theta_{sgf} || \theta_{hp}$ and Eq. (5) for θ and Eq. (3), respectively. Using the example described in Section 4.2, μ_{s_h} , σ_{s_h} and w_{gh} are updated in this equation. In this case, the penalty term prevents μ_{s_h} and $w_{gh, i}$, for all s_h , g_h , and i , from moving far away from μ'_{s_h} and $\mu'_{w_{gh, i}}$, respectively.

5 Experiments

In this section, we perform an experimental evaluation to investigate whether the symbol grounding functions and high-level planner are refined successfully by using the framework we proposed in the previous section. In Section 5.1, we focus on the evaluation for refining the symbol grounding functions only. Then, in Section 5.2, we evaluate the effect of jointly refining symbol grounding functions and the high-level planner.

5.1 Refinement of Symbol Grounding

We evaluate how the symbol grounding functions are refined by our framework to solve the mountain car problem. The experimental set up to implement the planner and our framework is the same as that in the example introduced in Section 3 and Section 4.

For the evaluation, we prepared three types of method:

Baseline: A hierarchical planner that uses the grounding functions and a high-level planner, which are manually designed. This planner is identical to that introduced in the example in Section 3.

NoPenalty: The framework that refines the symbol grounding functions without the penalty term in Eq. (17). In this method, the high-level planner is the same as that in Baseline.

Proposed: The framework that refines the symbol grounding functions with the penalty term. In this method, the high-level planner is the same as that in Baseline.

These methods were evaluated on the basis of two criteria: an average cumulative reward over episodes, and a parameter divergence. The former is to evaluate if the hierarchical planner produces a more appropriate plan by refining its modules, and the latter is to evaluate the interpretability of the refined modules. The parameter divergence represents how much the policy's parameters (μ_{s_h})³ refined by the framework differ from the initial parameters. In this paper, this divergence is measured by the Euclidean distance between the refined parameter (μ_{s_h}) and its initial parameter (μ'_{s_h}). Initial values for μ_{s_h} and σ_{s_h} are given, shown as "Init" in Table 3. μ_{s_h} is initialized with μ'_{s_h} , which is determined on the basis of the implementation of the functional in Eq. (6) (see Section 4.2), and σ_{s_h} is manually determined. We consider 50 episodes as one epoch and performed refinement over 2000 epochs.

The experimental results (shown in Figures 4 and 5) show that 1) refining the grounding functions improves the performance (average cumulative reward) of hierarchical planners, and 2) considering the penalty term keeps the refined parameters within a certain distance from the initial parameters. Regarding 1), Figure 4 shows the methods in which the grounding functions are refined (NoPenalty and Proposed) outperform Baseline. This result indicates the refinement for grounding functions successfully improves its performance. Regarding 2), Figure 5 shows that the parameter in NoPenalty moves away from the original parameter in refining, while in Proposed, the parameter stays close to the original one.

An example of the refined parameter for the grounding functions for Proposed is shown in Table 3, which indicates that the parameter is updated to achieve high-performance planning while staying close to the original parameter. In this example, the mean and standard deviation of "On_right_side_hill(Car)" is changed significantly through refinement. The mean for grounding On_right_side_hill(Car) is biased to a more negative position, and also flattened to make the car climb up the left side hill quickly (Figure 6). This change makes the symbol grounding function more flattened and considers the center position as

³We assume μ_{s_h} dominantly determines the behaviors of symbol grounding functions.

$$\theta_{sgf}||\theta_{hp} \leftarrow \theta_{sgf}||\theta_{hp} + \alpha V_{\pi_{\theta_{sgf}||\theta_{hp}}}(\tilde{s}_0) \underbrace{\{\nabla_{\theta_{sgf}||\theta_{hp}} \log P(\tilde{g}_0|\tilde{g}_{h,0}; \theta_{sgf}) P(\tilde{g}_{h,0}|\tilde{s}_{h,0}; \theta_{hp}) P(\tilde{s}_{h,0}|\tilde{s}_0; \theta_{sgf})\}}_{\text{reinforcement term}} + \underbrace{\nabla_{\theta_{sgf}||\theta_{hp}} \log P(\theta_{sgf}) P(\theta_{hp})}_{\text{penalty term}} \quad (17)$$




Figure 4: Learning curves for each methods. The vertical axis represents average cumulative rewards and the horizontal axis represents epochs (50 episodes for each epoch).

Figure 5: Parameter divergences. The vertical axis represents the average Euclidean distance and the horizontal axis represents learning epochs.

Table 3: Refined parameters.

| μ_{s_h} | Bottom_of_hills | At_top_of_right_side_hill | On_right_side_hill | On_left_side_hill |
|-------------|-----------------|---------------------------|--------------------|-------------------|
| Init | -0.5 | 0.6 | 0.2 | -1.1 |
| Refined | -0.5 | 0.46 | -0.39 | -1.1 |

| σ_{s_h} | Bottom_of_hills | At_top_of_right_side_hill | On_right_side_hill | On_left_side_hill |
|----------------|-----------------|---------------------------|--------------------|-------------------|
| Init | 0.4 | 0.1 | 0.4 | 0.3 |
| Refined | 0.4 | 0.12 | 1.42 | 0.11 |

“On_right_side_hill(Car).” The main interpretation of this result is that the symbol grounding function was refined to reduce the redundancy in high-level planning. In the original symbol grounding functions, the center position is grounded to “Bottom_of_hills (Car),” and the high-level planner makes a plan [Bottom_of_hills(Car) \rightarrow On_right_side_hill(Car) \rightarrow On_left_side_hill(Car) \rightarrow At_top_of_right_side_hill(Car)], which means “Starting at the bottom of hills, visit, in order, the right side hill, the left side hill, and the top of the right side hill.” However, this plan is redundant; the car does not need to visit the right side hill first. The refined symbol grounding function considers the center position as “Right_side_hill(Car),” and thus the high-level planner produces the plan [Bottom_of_hills(Car) \rightarrow On_right_side_hill(Car) \rightarrow On_left_side_hill(Car) \rightarrow At_top_of_right_side_hill(Car)], in which the redundancy is removed. It should also be noted that the order of the refined means is intuitively correct. For example, the value of $\mu_{\text{On_right_side_hill}}$ is higher than the value of $\mu_{\text{On_left_side_hill}}$ (i.e., $\mu_{\text{On_right_side_hill}}$ means the place on more right-side than $\mu_{\text{On_left_side_hill}}$). It cannot be seen in the Baseline and NoPenalty cases. This result supports the fact that our framework refines the modules by maintaining their interpretability.

5.2 Joint Refinement of Symbol Grounding and High-Level Planning

In this section, we refine both the symbol grounding functions and the high-level planner. The setup of the hierarchical planner and the problem are the same as those of the previous section, except for the knowledge base for the high-level planner. We removed “Opr.2” (as shown in Table 1) and used this degraded version as the knowledge base for the experiment.




Figure 6: Example of refining result symbol grounding function for *On_right_side_hill*.

Figure 7: Learning curve.

This degradation makes a space for refining the knowledge base for the high-level planner. In addition, we put a small coefficient of the penalty term for w_{gh} , because we found that considering this term too much makes the refinement worse in a preliminary experiment. As long as the results of the symbol grounding functions are interpretable, the result of the high-level planner is interpretable as well. w_{gh} is initialized with $w'_{gh,i}$, which is determined by (i.e., Algorithm 1) where we set -1.3 as val_{in} , and -0.02 as val_{min} . The resulting $w'_{gh,i}$ is shown as “Init” in Table 4.

We prepared three types of methods:

NoRefining: A hierarchical planner with the degraded version of the knowledge base for high-level planner. The knowledge base for the symbol grounding function is the same to that shown in Table 2.

RefiningHP: The framework that refines the high-level planner only. In this method, symbol grounding functions are the same as those in NoRefining.

RefiningHPSGF: The framework that refines both symbol grounding functions and the high-level planner.

From the experimental result (Figure 7), we can confirm that our framework successfully refines both symbol grounding functions and the high-level planner, from the viewpoint of performance. RefiningHP outperforms NoRefining, and RefiningHPSGF outperforms the other methods.

Table 4 provides an example of how the high-level planner was refined. It indicates that the dropped knowledge (i.e., Opr. 2) was successfully acquired in refinement, and knowledge is discovered that makes high-level planning more efficient. Considering the form of Eq. (12), the operator, which corresponds to the element of a weight with a higher value, contributes more to high-level planning. Therefore, these corresponding operators are worthwhile as knowledge for high-level planning. In Table 4, the refined weight of the operator (preconditions=On_right_side_hill, effects=On_left_side_hill) is higher than those of other operators in which the precondition contains On_right_side_hill. This operator was once initially removed and later acquired by the refinement. Similarly, the operator (preconditions=Bottom_of_hills, effects=On_left_side_hill), which is not shown in Table 1, was newly acquired.

Table 4: Example of high-level planner improvement. Refined weights w_{gh} are shown for each precondition (column) and effect (row). Initial weights are shown in parentheses.

| Refined (Init) | Bottom_of_hills | At_top_of_right_side_hill | On_right_side_hill | On_left_side_hill |
|---------------------------|--------------------|---------------------------|--------------------|--------------------|
| Bottom_of_hills | -5.88 (-1.3) | -6.34 (-1.3) | -3.15 (-1.3) | -6.65 (-1.3) |
| At_top_of_right_side_hill | -9.04 (-1.3) | -9.75 (-1.3) | -4.76 (-1.3) | 2.5 (-0.02) |
| On_right_side_hill | -0.98 (-0.02) | 1 (-1.3) | -2.03 (-1.3) | -1.34 (-1.3) |
| On_left_side_hill | 0.85 (-1.3) | -2.12 (-1.3) | 1.74 (-1.3) | -11.71 (-1.3) |

6 Conclusion

In this paper, we proposed a framework that refines manually-designed symbol grounding functions and a high-level planner. Our framework refines these modules with policy gradients. Unlike standard policy gradient implementations, our framework additionally considers the penalty term to keep parameters close to the prior parameter derived from manually-designed modules. Experimental results showed that our framework successfully refined the parameters for the modules; it improves the performance (cumulative reward) of the hierarchical planner, and keeps the parameters close to those derived from the manually-designed modules.

One of the limitations of our framework is that it deals only with predefined symbols (such “Bottom_of_hills”), and it does not discover new symbols. We plan to address this drawback in future work. We also plan to evaluate our framework in a more complex domain where primitive actions and states are high-dimensional, and the knowledge base is represented in a more complex description (e.g., precondition contains multiple states).

References

- [Camacho and Alba, 2013] Eduardo F Camacho and Carlos Bordons Alba. *Model predictive control*. Springer Science & Business Media, 2013.
- [Cambon *et al.*, 2009] Stéphane Cambon, Rachid Alami, and Fabien Gravot. A hybrid approach to intricate motion, manipulation and task planning. *The International Journal of Robotics Research*, 28(1):104–126, 2009.
- [Choi and Amir, 2009] Jaesik Choi and Eyal Amir. Combining planning and motion planning. In *Proc. of ICRA-09*, pages 238–244. IEEE, 2009.
- [Das *et al.*, 2017] Abhishek Das, Satwik Kottur, José M. F. Moura, Stefan Lee, and Dhruv Batra. Learning cooperative visual dialog agents with deep reinforcement learning. arXiv:1703.06585, 2017.
- [Dornhege *et al.*, 2009] Christian Dornhege, Marc Gissler, Matthias Teschner, and Bernhard Nebel. Integrating symbolic and geometric planning for mobile manipulation. In *Proc. of SSR-09*, pages 1–6. IEEE, 2009.
- [Fdez-Olivares *et al.*, 2011] Juan Fdez-Olivares, Luis Castillo, Juan A C  zar, and Oscar Garc  a P  rez. Supporting clinical processes and decisions by hierarchical planning and scheduling. *Computational Intelligence*, 27(1):103–122, 2011.
- [Kaelbling and Lozano-P  rez, 2011] Leslie Pack Kaelbling and Tom  s Lozano-P  rez. Hierarchical task and motion planning in the now. In *Proc. of ICRA-11*, pages 1470–1477. IEEE, 2011.
- [Kakade, 2002] Sham M Kakade. A natural policy gradient. In *Proc. of NIPS-02*, pages 1531–1538, 2002.
- [Konidaris *et al.*, 2014] George Konidaris, Leslie Pack Kaelbling, and Tomas Lozano-P  rez. Constructing symbolic representations for high-level planning. In *Proc. of AAAI-14*, 2014.
- [Konidaris *et al.*, 2015] George Konidaris, Leslie Pack Kaelbling, and Tomas Lozano-P  rez. Symbol acquisition for probabilistic high-level planning. In *Proc. of IJCAI-15*, 2015.
- [Konidaris, 2016] George Konidaris. Constructing abstraction hierarchies using a skill-symbol loop. In *Proc. of IJCAI-16*, 2016.
- [Malcolm and Smithers, 1990] Chris Malcolm and Tim Smithers. Symbol grounding via a hybrid architecture in an autonomous assembly system. *Robotics and Autonomous Systems*, 6(1-2):123–144, 1990.
- [McDermott *et al.*, 1998] Drew McDermott, Malik Ghallab, Adele Howe, Craig Knoblock, Ashwin Ram, Manuela Veloso, Daniel Weld, and David Wilkins. PDDL-the planning domain definition language. 1998.
- [Moore, 1991] Andrew Moore. Efficient memory-based learning for robot control. March 1991.
- [Nilsson, 1984] Nils J Nilsson. Shakey the robot. Technical report, SRI INTERNATIONAL MENLO PARK CA, 1984.
- [  zdamar *et al.*, 1998] Linet   zdamar, M Ali Bozoyel, and S Ilker Birbil. A hierarchical decision support system for production planning (with case study). *European Journal of Operational Research*, 104(3):403–422, 1998.
- [Schulman *et al.*, 2015] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *Proc. of ICML-15*, pages 1889–1897, 2015.
- [Silver *et al.*, 2014] David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *Proc. of ICML-14*, 2014.
- [Silver *et al.*, 2016] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of Go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- [Sutton *et al.*, 1999] Richard S Sutton, Doina Precup, and Satinder Singh. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2):181–211, 1999.
- [Sutton *et al.*, 2000] Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Proc. of NIPS-00*, pages 1057–1063, 2000.
- [Williams, 1992] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. In *Reinforcement Learning*, pages 5–32. Springer, 1992.
- [Wolfe *et al.*, 2010] Jason Andrew Wolfe, Bhaskara Marthi, and Stuart J Russell. Combined task and motion planning for mobile manipulation. In *Proc. of ICAPS-10*, 2010.