Heuristically Adaptive Policy Reuse in Reinforcement Learning

Han-Chao Wang, Tianpei Yang, Jianye Hao*
College of Intelligence and Computing, Tianjin University
Tianjin, China
ares1899@126.com,tpyang@tju.edu.cn,jianye.hao@tju.edu.cn

ABSTRACT

Transfer learning can significantly improve the reinforcement learning process by leveraging prior knowledge from past learned tasks. However, how to select an optimal source task for reuse to improve a reinforcement learning agent is still challenging. In this paper, we propose a new Policy Reuse framework called Heuristically Adaptive Policy Reuse (HAPR) that facilitates efficient reuse of source policies, which is stored in a given Policy Library, by rapidly selecting the most appropriate policy only with its useful part. For the agent to reuse in successive tasks, HAPR is also capable of rebuilding the Policy Library to provide representative policies, whose quality is guaranteed by using KL-divergence to measure the irrelevance between policies. Our extensive experiments based on a grid-world domain show the efficiency and robustness of our method, compared with the state-of-the-art policy reuse approaches.

KEYWORDS

Policy Reuse; Transfer Learning; Reinforcement Learning

1 INTRODUCTION

Reinforcement learning (RL) [19] is a proverbial framework for an agent to learn and optimize a policy through the interactions with the environment and feedback rewards. Most RL methods are faced with sample efficiency problems, which makes it difficult to learn from scratch, especially in solving complex tasks. Transfer learning (TL) [20] can facilitate RL to improve learning efficiency by transferring past learned knowledge to a new task, which usually consists of three phases: select useful knowledge to transfer, find a suitable way to transfer and last, execute the transfer process. If the transferred knowledge which is incompatible with the target environment is selected, negative transfer occurs [12]. It is challenging to predict whether source information is useful in advance.

Policy Reuse is useful to accelerate RL process by reusing past similar learned policies. Existing researches on policy reuse include reusing expert suggestions [3, 18], defining the policy similarities in a reward shaping manner [2], modeling policy selection as Bayesian optimization problems [13], and transferring the experience instances of a source task to the target task by reusing these instances to estimate the reward function [9]. However, these approaches require more extra knowledge for effective transfer. Fernández and Veloso [5] proposed Policy Reuse Q-learning (PRQL) and Policy Library rebuilding through Policy Reuse (PLPR) where a learning agent is equipped with a library of previous policies to facilitate exploration, as they enable the agent to collect relevant information quickly to accelerate learning. However, PRQL converges to a sub-optimal policy in some conditions since negative transfer exists, and PLPR builds a Policy Library without a clear theoretical guarantee. Li and Zhang [10] proposed an Optimal Policy Selection TL method (OPS-TL) to select a suitable policy using a multi-armed bandit (MAB) method UCB1 [8] during the online learning. However, OPS-TL needs more performance feedback to evaluate sources to selection, which takes time to lock a known suitable policy. And it also requires a manual setting of rate to learn independently.

To address the above problems, we propose a Heuristically Adaptive Policy Reuse (HAPR) framework as our contribution. This framework consists of two parts: the primary part of HAPR for Transfer Learning (HAPR-TL) and a Policy Library rebuilding part using KL-divergence (PLKL) as secondary. HAPR-TL reuses a policy all out during learning until a state called *subgoal* is reached, where the policy is selected by evading other inappropriate source policies. A subgoal of a source policy is a special state defined by using a state-value function, which is also used to filter out bad policies or parts. The idea of subgoal is from the concept proposed by Ruby and Kibler [15]. In order to make HAPR-TL persistent and become lifelong learning, the secondary PLKL method rebuilds the library once HAPR-TL finishes. It uses KL-divergence [7] to measure the distance between two policies and determine whether to add a new policy into the policy library or not. Our experimental results on the grid-world domain show our method outperforms state-of-the-art policy reuse approaches, as HAPR adaptively reuses the useful knowledge from a policy library that is equipped with various policies without redundancy.

The main contributions of this paper can be summarized as:

(1) To reuse policy in a state-level and to avoid reuse from the irrelevant part of the source policy, we draw lessons from the subgoal method and set the subgoal as a critical state to stop policy reusing.
(2) To avoid negative transfer with reusing known unsuitable source policy, we give up the UCB1 exploration. We choose the successful ratio to the goal of tasks as the evaluation of source policies and remove the poor policy from the alternative Policy Library in each episode.
(3) To make the Policy Library more representative, we use KL-divergence to rebuild the Policy Library.

The remainder of this paper is organized as follows. Section 2 introduces the background of our approach including problem formulation in 2.1 and related works in 2.2 and 2.3. Section 3 presents our approach in two subsections: Policy Reuse in 3.1 and Policy Library rebuilding in 3.2. Section 4 gives the experimental results compared with state-of-the-art methods. Section 5 concludes our approach and draws the future work.

*Corresponding author: Jianye Hao

2 BACKGROUND

2.1 Policy Reuse Problem

RL problems are usually formulated as Markov Decision Processes (MDPs), which is defined in a 5-tuple $< S, A, T, R, \gamma >$. S is a finite set of states. A is a finite set of actions. T is a stochastic state transition function ($T = S \times A \times S \to [0, 1]$). R is a stochastic reward function ($R = S \times A \times S \to \mathbb{R}$). And γ is a discounted factor ($\gamma \in [0, 1]$). A policy π is defined as a function that specifies an appropriate action $a = \pi(s)$ for each state s. An agent following π will get a discounted total reward $W_\pi = \sum_{h=1}^{H} \gamma^{h-1} R(s_h, a_h)$, with reward $R(s_h, a_h)$ feedback in step h. The solution of an MDP is to find an optimal policy $\pi^* = \arg\max_\pi \mathbb{E}[W_\pi]$ to maximize the expected value of W_π. In practice, the mean of sampled reward $\overline{W}_\pi$ is also an evaluating indicator for the algorithm performance.

To formally describe the *policy reuse problem*, a domain $\mathcal{D}$ is defined as a sub-tuple $< S, A, T >$ of MDPs. And a task Ω is defined as a tuple $< \mathcal{D}, R_\Omega >$ with domain $\mathcal{D}$ and R of MDPs. A policy library L is a set of source policies $\pi_1, \pi_2, \ldots, \pi_n$, where policy π_i is a trained policy of task Ω_i. With different task intervals in the same domain $\mathcal{D}$, the *policy reuse problem* is to find a way to learn the optimal policy π^*_Ω of every new task Ω by reusing source policy selected from a given policy library L.

2.2 Policy Reuse Methods

Compositional Q-learning [16] first prompted the problem of learning multiple tasks in the same domain by TL. In addition, a different Policy Reuse context for the lifelong autonomous agent was described by Rosman *et al.* [13]. PRQL [5] optimize gave a classic Policy Reuse framework to solve this problem. PRQL reuses policy, which is selected from the Policy Library with its own policy being trained following the soft-max (Boltzmann distribution) strategy, as the training policy in a certain probability. A current method OPS-TL [10] optimizes policy selection method of PRQL by using MAB method of UCB1 [8]. However, the exploration rate of OPS-TL is also increased exponentially without effective reuse in a training episode. And with no contribution, the UCB1 method sometimes selects a known poorer policy for exploration.

Reuse with exploration leads to that PRQL and OPS-TL can't quickly learn a task similar to the source task. And the reserved selection makes PRQL and OPS-TL cannot get rid of the condition of reusing unsuitable policy quickly, which leads to negative transfer at the beginning. By contrast, Our Policy Reuse method HAPR-TL will lock the selection of the most suitable policy and fully reuse the useful part of the policy.

2.3 Policy Library Rebuilding Methods

Policy Library rebuilding method provides source policies to serve future Transfer Learning and responds to the domain structure to a certain extent.

PLPR proposed by Fernández and Veloso is the only Policy Library rebuilding method that is associated with Policy Reuse. However, PLPR lacks intuitive explanation. For library rebuilding and allowing the mechanics of the environment to be learned, a parallel learning method is proposed by Ollington and Vamplew [11]. However, problems don't arise just right together in reality.

In addition, the method of Earth Mover's Distance (EMD) [14] using Wasserstein Metric can also be used to present the dissimilarity between policies as an improved method of PLPR. The recent work on representing similarity between distributions was proposed by Song *et al.*, which defines a distance between MDPs using EMD [17]. However, computing EMD needs a high time complexity.

Our Policy Library rebuilding method PLKL optimizes PLPR with the theoretical support of KL-divergence [7]. So that PLKL can further solve sequential tasks by using previous source policies as PLPR does. PLKL computes the KL-divergence in both directions of each two policies as the criteria for rebuilding the policy library simply and sufficiently.

3 APPROACH

Our approach addresses the *policy reuse problem*. It starts with a Policy Library L, and it uses a policy reuse method to learn an optimal policy of a target task with a source policy, which is selected by a policy selection method from L. Our approach has also rebuilt L for the next task to learn. In this section, we give our policy reuse framework HAPR to solve the Policy Reuse problem, which concludes its main part HAPR-TL for policy reuse, with a library rebuilding method PLKL to assist it, as shown in Algorithm 1.

Algorithm 1 HAPR

Require: Policy Library L

```
loop
    get a new task Ω
    π_Ω ← HAPR-TL(Ω, L)
    L ← PLKL(L, π_Ω)
end loop
```

In Algorithm 1, given a policy library L and a new task Ω, HAPR-TL learns an optimal policy of Ω by fully reusing policies selected from L. And PLKL rebuilds the Policy Library L once obtaining the learned policy. The details of each part are shown in the following section 3.1, and section 3.2, which are two main components of our transfer learning framework.

3.1 Policy Reuse with Selection for TL

Considering that the agent cannot foresee whether the knowledge of source policy is suitable for the target task to transfer, our approach focuses on the way to select the useful part from given source policies for reuse. In this section, we introduce our HAPR-TL method in two parts following: Policy Selection and Policy Reuse.

Policy Selection in HAPR-TL. An intention of Policy Reuse is to reuse useful parts of source policies not to well-train before. In our approach, the source policy selection method has three purposes.

(1) To quickly lock the most suitable policy for the target task.
(2) To get the useful part of a selected source policy to reuse.
(3) To stop reuse once the target policy performs as good as the source policy.

In our algorithms: S stands for the state set containing all the available states in our domain; L stands for the Policy Library; and L' stands for the policy set for selecting source policies. And $|S|$, $|L|$

Algorithm 2 HAPR-TL(; L)

Require: States Set S; target task with its goal state s_G; source policies in Policy Library L with the Q-function, the C-functions and the subgoal set for each policy

1: Initialize Target Policy : $\forall Q(s,a) = 0, \forall C(s) = 0$
» ¼; $L' \leftarrow L$; $s \leftarrow s_G$; $n_i \leftarrow 1, p_i \leftarrow 0.5 \ \forall i = 1..|L|$
2: for n = 1::N do
3: if $s_G <$ or $\forall s <$ for subgoals s of any policy in L' then
4: if n > 1 then
5: $p_k \leftarrow \frac{p_k \cdot n_k}{n_k + 1}$
6: end if
7: k the next index i of $\in L'$
8: else
9: $p_k \leftarrow \frac{p_k \cdot n_k + 1}{n_k + 1}$
10: for $\forall \in L'$ do
11: if $2 \cdot p_i < \max_{j \in L'} p_j$ then
12: $L' \leftarrow L' - \{i\}$
13: else
14: $g(i) \leftarrow \sum_{s \in »s;s_G¼} \min\{0, C_i(s) - \&\}$
15: end if
16: end for
17: $k \leftarrow \arg\max_i \{g(i) \mid i \in L'\}$
18: end if
19: $n_k \leftarrow n_k + 1$
20: Get uniform random Initial State $s_0 \in S$
21: Get the subgoal of $_k$ by and the performance of .
22: [,] _reuse($_k$; ; s_0; s; s_G;)
23: end for
24: return

and $|L'|$ are the numbers of elements in those sets. & is a mean-value threshold of the C-function: $\& := \frac{1}{|S|} \cdot$.

Algorithm 2 takes N episodes to train the policy for. In each episode, our selection method first selects a source policy to reuse (Line: 3-18) and then revise the state of subgoal (Line: 21). After that is our Policy Reuse method (Line: 22; shown in Algorithm 3).

For each episode, is the sampled state trajectory, which records all the states passed in the previous episode in order. Only if reached the goal of the target task and also reaches at least one subgoal of any source policy in L' (Line: 8), our policy selection mechanism will really be implemented (Line: 10-17). Otherwise, it will select a policy and a subgoal in turn (Line: 3,7).

A success rate p to the goal of a source policy is adjusted in every episode to indicate the reliability of the policy (Line: 5, 9). Our selection method stops reusing of poor policy (with a low p) by removing the policy from the library L' for selecting (Line 11-12).

For each source policy $_{src}$, we first require a subgoal set to select a subgoal for our Policy Reuse method (Line: 21, 22). The subgoal set includes the subgoals and the goal of $_{src}$. If some states in the subgoal set exist in the trajectory, our method will revise the current subgoal to one of those states last appeared in. Otherwise, it will select a subgoal of $_{src}$ in turn. Furthermore, if the target policy has a higher mean value of total reward $\bar{W}$ than any source policy $_{src}$ in a testing batch, it will degenerate to the initial state s_0 of the current episode.

In our approach, we propose C function to present the importance of states. If a source policy only contains Q-function, we will add C-function to the policy. The values of C can be figured out by generating a series of trajectories guided by Q. C updates every step in a sampled trajectory. In a step from state s' to s:

$$C(s') \leftarrow (1 -) \cdot C(s') + \cdot C(s) \quad (1)$$

where C is updated in the form of the linear combination of $C(s')$ and $C(s)$. Stepsize shares the same value with the stepsize of iteration in our approach. Finally, C needs to be normalized. We use C-function for both comparing the source policies and obtaining the subgoal set.

For the trajectory reaching the goal and the subgoal in the previous episode, we calculate the cumulative value of C over the threshold & for every source policy in L' to compare and select an outstanding source policy $_k$ (Line: 14, 17):

$$k \leftarrow \arg\max_i \sum_{s \in »s;s_G¼} \min\{0, C_i(s) - \&\} \quad (2)$$

where $»s;s_G¼$ is the set including all the state from s to s_G that the trajectory passes through.

Sharing the same idea with subgoal method [5], we use C-function of each source policy to figure out its subgoal set as Algorithm2 requires. In the h^{th} step of a sampled trajectory based on a source policy, that the state s_h is a subgoal of the policy subjects to:

$$C(s_h) > \max(, C(s_h) +) \quad (3)$$

where

$$C(s_h) = C(s_h) - C(s_{h-1}) \quad (4)$$

In Formula 3 and 4, C and are two positive thresholds, which indicates that $C(s_h) > 0$ for the subgoal s_h. For each source policy, our approach only adopts its subgoals in a small number, which number can be effectively controlled by adjusting these two thresholds.

The C-function can also be used to examine how frequently the training has worked in a state, and also to figure out the KL-divergence between policies.

The HAPR-TL method finishes with returning the policy derived by Q with C included. Our Selection method uses the notion of subgoal and C-function to avoid negative transfer from source policies. This method is constructed to match the following Policy Reuse method.

Policy Reuse of HAPR-TL. As we hand over the duty of sifting the useful part of source policy to the Policy Selection method, in Algorithm 2, we character our Policy Reuse method to completely reuse the given policy. The subgoal selected in Algorithm 2 is used as a signal in our Policy Reuse method shown in Algorithm 3. Our method keeps reusing the source policy until a given subgoal is reached and then turns to learn with its own policy.

To reuse policy fully in H steps, we keep reusing policy until a provided subgoal reached (Line: 3-4). After the step arriving in an episode: if we have arrived the goal of the target task in the previous episode, we will use the greedy method to correct our target policy; otherwise we use the random policy to enhance our exploration (Line: 5-10).

A new state trajectory $_{neo}$ consists of the initial state s_0 and all the state arrived after an action in each step (Line: 1,12). When

Algorithm 3 _reuse ($_{src}$; ; s_0; s; s_G;)

Require: Source Policy $_{src}$; Target Policy with its Q-value function and C-value function; Initial State $s_0 \in S$; State of subgoals s; Goal States s_G; last State Trajectory

1: Initialize $_{neo}$ » s_0¼
2: for h = 1::H do
3: if s < $_{neo}$ then
4: $a_{h\ 1}$ $_{src}$¹$s_{h\ 1}$º
5: else
6: if $s_G \in$ then
7: $a_{h\ 1}$ _greedy¹ ¹$s_{h\ 1}$ºº
8: else
9: Get uniform random $a_{h\ 1} \in A$¹$s_{h\ 1}$º
10: end if
11: end if
12: Get state s_h after the action $a_{h\ 1}$ and put s_h in $_{neo}$.
13: Update Q:
Q¹$s_{h\ 1}$; $a_{h\ 1}$º
¹1 º Q¹$s_{h\ 1}$; $a_{h\ 1}$º + ¹r_h + max$_a$¹Q¹s_h; aºº
14: if $s_{h\ 1} <$ $_{neo}$ then
15: C^{0}¹s_hº ¹ 1 º C¹s_hº + C¹$s_{h\ 1}$º + 1
16: end if
17: if $s_h = s_G$ then
18: Update C: C C^0
19: exit
20: end if
21: end for
22: Normalize C with sum of 1
23: return ; $_{neo}$

getting a new state s after an action from s^0, Q-function is updated. The iteration of C-function works at the same time with Q, but it is just temporarily updated:

$$C^0(s^0) \leftarrow (1-\alpha)\, C(s^0) + C(s^0) + 1 \tag{5}$$

Formula 5 is slightly different from Formula 1 because $C(s^0)$ is not really updated. Considering to update more unbiased, $C(s^0)$ is only temporarily updated in the first time reaching a state during the current episode. And C is really updated only if the current episode finally arrives to the goal state s_G (Line: 14-18). Algorithm 3 finished if s_G is arrived or time step H is up. And C-function also needs to be normalized here (Line: 22).

As our Policy Reuse method maximizes the reuse of the policy part selected, the method will have a good performance at the very beginning if the part reused is useful, or it will feedback a bad performance in time and the algorithm will avoid selecting the policy having a negative correlation with the target task. And it also can prevent the unbalanced sampling of trajectories in training, so that the new C generated can be used as C-function within the source policy for the PLKL method in the next subsection.

3.2 Policy Library rebuilt with KL-divergence

Follow the previous work settings [5], the policy reuse problem often considers the situation that an agent is equipped with a policy library, although HAPR-TL can also learn with itself. In addition, a policy library can indirectly represent the dynamics of the environment, which also helps the learning process of the agent.

In this section, we give a policy library rebuilding method to ensure the independence among source policies. First, to ensure that each policy in the Policy Library is unique and independent with each other, we use KL-divergence (relative entropy) [7] as a simple criterion to measure the dissimilarity between two policies. Our PLKL method takes the advantages of the KL-divergence of the C-functions of a source policy and the target one in both directions shown in Algorithm 4.

Algorithm 4 PLKL(L;)

Require: Policy Library L with C-value functions; Target Policy with its C-value function

1: for $_i \in$ L do
2: D_{KL} 0; $D_{KL\ inv}$ 0
3: for 8s 2 S do
4: if C ¹s º > 0 and C_i ¹s º > 0 then
5: D_{KL} $D_{KL} + C_i$ ¹s º log¹$\frac{C_i\ ¹s º}{C\ ¹s º}$º
6: $D_{KL\ inv}$ $D_{KL\ inv}$ + C ¹s º log¹$\frac{C\ ¹s º}{C_i\ ¹s º}$º
7: end if
8: end for
9: if $D_{KL} <$ then
10: return L
11: else if $D_{KL\ inv} <$ then
12: L L f $_i$ g
13: end if
14: end for
15: return ¹L [f g º

We calculate the KL-divergence between the source policy and the target policy to measure the former can be replaced by the latter, shown in Formula 6:

$$D_{KL} = \sum_{8s\in S,\, 8C(s)>0} {}^{1}C_{src}(s)\ \log{}^{1}\frac{C_{src}(s)}{C_{tar}(s)}{}^{oo} \tag{6}$$

where the source policy's C-function C_{src} and the target policy's C-function C_{tar} can be considered as normalized distributions of the importance of states in their respective tasks. A large value of D_{KL} in Formula 6 indicates that the original policy is not similar with the target policy. In Algorithm 4, if every D_{KL} exceeds a threshold $_D$ = (2 <¹ 0; 1¼, the new policy will be joined in the Policy Library L (Line 15). We refer the PLKL only considering the uni-direction KL-divergence (without Line: 11,12) as the uni-direction PLKL (uni-KL) method.

Similarly, the inverse KL-divergence can find out whether the target policy can replace a source policy:

$$D_{KL\ inv} = \sum_{8s\in S,\, 8C(s)>0} {}^{1}C_{tar}(s)\ \log{}^{1}\frac{C_{tar}(s)}{C_{src}(s)}{}^{oo} \tag{7}$$

A small value of $D_{KL\ inv}$ value in Formula 7 indicates that the target policy can replace the source policy. In Algorithm 4, besides joining the target policy in the Policy Library, when the value of $D_{KL\ inv}$ is lower than $_D$ while D_{KL} is not, the new policy is decided to replace the source policy in (Line: 11,12,15). We refer the whole

PLKL method as the bi-direction PLKL (bi-KL) method. According to the asymmetry of KL-divergence, this method can remove similar policies from the source Policy Library.

Our PLKL method ensures that each policy is independence from each other with a theoretical guarantee, as KL-divergence can describe the difference from one distribution to another. KL-divergence is not the only option for the requirement, Bhattacharyya distance [1] or JS-divergence [4] can also be used here.

4 EXPERIMENTS

In this section, we provide experimental results on a grid-world navigation domain. We first give some heat-maps to show the feasibility of the C-function in our results. Then we verify that our Policy Reuse method HAPR can transfer quickly and avoid negative transfer, compared with the related methods including ε-greedy Q-learning, PRQL [5] and OPS-TL [10]. We also compare our PLKL with libraries rebuilt by PLPR [5].

4.1 Experimental Setting

Our experiments use a 24×21 grid-world navigation domain with 50 sequential tasks, which are represented by their goals in Figure 1.

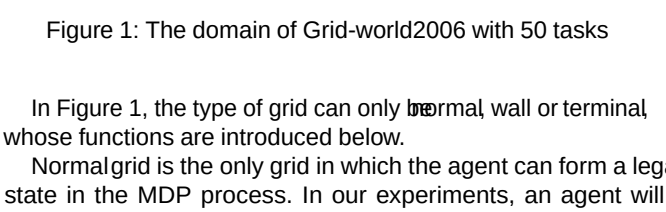

Figure 1: The domain of Grid-world2006 with 50 tasks

In Figure 1, the type of grid can only be normal, wall or terminal, whose functions are introduced below.

Normal grid is the only grid in which the agent can form a legal state in the MDP process. In our experiments, an agent will be randomly generated in a normal grid to start a navigation episode with a number of steps. In each step, the agent will choose one direction from east, north, west and south with a distance to move as an action. After an action, the agent will access to a new grid. If the new grid is the wall, the agent will be transferred back to the last position and waste one step. If the new grid is the terminal, a reward of arriving the goal state (only arriving terminal has reward) will be received and the current episode should be terminated. The terminal must be fixed in our goal-oriented task.

For ease of description, we named every grid by its column and row start from "Grid(0,0)" on the top-left, such as "Grid(3,2)" refers to the terminal grid of task 2 in Figure 1. Without loss of generality, we set the length of grids' edge to 1. And we represent the position of agent within a two-dimensional continuous coordinate as (x_a, y_a), where $x_a \in (0, 24)$ and $y \in (0, 21)$. The agent is in grid (x, y) when $(x_a, y_a) = (x, y)$, where $x \in \mathbb{N}(0, 24)$ and $y \in \mathbb{N}(0, 21)$. Each action can change one coordinate of the agent's position in length 1 in a direction. To east increase the "x_a"; to north decrease the "y_a"; to west decrease the "x_a" and to south increase the "y_a". The actual arrival position is affected by an error following a uniform distribution in a range of $(-0.20, +0.20)$.

This environment has $|S| = 301$ accessible normal states. We set $\alpha = 0.05$, $\gamma = 0.95$ for Q-function & C-function update for all method in experiment. And $\varepsilon = 0.90$ as the exploit rate of the ε-greedy method. The parameters of the comparison methods including PRQL [5] and OPS-TL [10] are consistent with the best in those methods. Each method will be trained in $N = 4000$ episodes and at most $H = 100$ steps within an episode. And for each $N^{10} = 100$ episodes, there is a set of test episodes to evaluate the performance of the target policy trained.

4.2 Experiment Results

In our experiments, the well-trained policies of tasks $_1$, $_2$, $_3$ and $_4$, whose goals are shown in Figure 1, are chosen into an initial source Policy Library $\Lambda_{init} = \{ _1; _2; _3; _4\}$.

Feasibility of C-function. In the first experiment, we choose tasks $_{46}$ and $_{29}$ shown in Figure 1 as target tasks. We compare the C-function of them with the tasks in Λ_{init}.

It is intuitive to see that the tasks corresponding to the same room ($_{46}$ and $_2$) in Figure 1 are similar. And the heat-maps in Figure 2 shows that such similar tasks have great similarity with their C-functions, which is normalized by the sum of 1 in this experiment. Figure 1 also shows that $_{29}$ didn't have any similar task from Λ_{init}. This result manifests that the C-function can be used as a representative feature of a task. According to this result, we respectively choose $_{46}$ and $_{29}$ as the target tasks in the second experiment and the third experiment.

In addition, we get the subgoals of each source policy in preparation for the next experiments. Instead of setting values for α and β in Formula 3, we sort the states directly according to the form of Formula 3 and choose the first b ($b = 2$) states as subgoals for each source policy $_{src}$. First, we sample several trajectories according to $_{src}$ and figure out $C(s)$ and $C(s')$ for each state s and the next s' in every trajectory. Then, we find out every subgoal s with $C(s) > 0$ and $C(s') - C(s) \leq 0$. We rank them according to their $C(s)$. If their number is more than b, we take the first b states of them as subgoals in $_{src}$'s subgoal set. Otherwise, until the number of subgoals reaches b, we will keep picking up the state with the highest value of $\frac{C(s)}{C(s')}$ from the rest with $C(s) > 0$.

In our experiment, the subgoal sets for the source policies in Policy Library Λ_{init} are shown in Table 1. For each task, the states in subgoal set consists of two subgoals and the goal of the task.

Figure 3: Average discounted rewards with suitable τ_{src}

1000 episodes to a value more than 0.4. Compared with PRQL and OPS-TL, HAPR has the advantage at "Jumpstart", "Asymptotic Performance", "Time to Threshold" and other evaluate metrics proposed by Taylor and Stone [20]. Our method learns quickly mainly because we fully reuse the source policy with a quick selection. However, the other method have slow selection, and their reuse rate updates as ' ' with $= 0.95$ in every step, which leads to a low learning rate' even in the tenth step ($0.95^{10} < 0.60$).

Figure 2: C-functions of $_1$; $_2$; $_3$; $_4$; $_{46}$ and $_{29}$

Table 1: Subgoal and goal for each task in L_{init}

Task	subgoals S_1	subgoals S_2	goals S_G
1	Grid(13,5)	Grid(16,2)	Grid(18,1)
2	Grid(6,6)	Grid(5,5)	Grid(3,2)
3	Grid(5,15)	Grid(2,15)	Grid(3,18)
4	Grid(17,15)	Grid(18,18)	Grid(20,18)

HAPR-TL with a similar task. In the second experiment, we choose the task $_{46}$ shown in Figure 1 as the target task. It obviously has a similar task $_2$ in the Policy Library L.

Figure 3 shows the learning curves of HAPR-TL in our approach, OPS-TL, PRQL and -greedy QL top-down when solving task $_{46}$. The learning curve is generated by the average discounted reward $\overline{W}$ of each method's on-policy testing, which executed 10 times after every 100 episodes from 100 to 4000. Error bars of standard deviations is shrinking to half.

In Figure 3, the average reward W of HAPR-TL is greater than 0.3 at starting with the first 100 episodes converges quickly in about

Figure 4: Frequency of reuse from each source task to $_{46}$

Figure 4 shows frequency curves of the reuse rate of each source policy in the Policy Library L, where the curve of reusing policy of $_2$ is obviously higher than other curves in all three algorithms as HAPR-TL, OPS-TL, and PRQL. Error bars represent standard deviations.

In Figure 4, HAPR-TL quickly locked the source policy at the beginning with the highest rate near 100%. It shows that our selection method is effective to select the right policy. In HAPR-TL, the exploitation rate for reusing policy of $_2$ did not go down as in

PRQL or in OPS-TL, because our method has defined the subgoal for policy reuse problem and actually learns independently when the subgoal degenerate to the initial state in an episode.

HAPR-TL without any similar task. In the third experiment, we choose task $_{29}$ as the target task, since it has no similar task in the Policy Library mentioned in the first experiment.

Figure 5: Average discounted rewards with unsuitable $_{src}$

Figure 5 shows the learning curve of HAPR-TL, ϵ-greedy QL, OPS-TL, and PRQL of $_{29}$ in top-down order. The curve in Figure 5 is generated in the same way as in Figure 3.

In Figure 5, HAPR-TL can be found to basically exceed ϵ-greedy QL, while OPS-TL is just close to ϵ-greedy QL. PRQL performs the worst. Our method performs well at the beginning mainly because it gets rid of the bad part of policies to reuse since all source policies are not able to completely reuse.

Figure 6: Frequency of reuse from each source task to $_{29}$

Figure 6 shows frequency curves of each source policy for learning $_{29}$, in the same way as Figure 4. The curve of reusing policy of $_{4}$ is obviously higher than other curves in HAPR-TL and OPS-TL.

In Figure 6, without a suitable source policy, reusing whole source policy quickly abandoned as the frequency curves of reusing each policy bifurcates in about 300 episodes. However, OPS-TL needs 500 episodes to start this. Our method is better than OPS-TL and PRQL in the absence of the source policy with its stability since we still reuse part of policy $_{4}$.

Therefore, experiments 2 and 3 empirically demonstrate that HAPR-TL significantly lowers the sample complexity of reaching convergence.

Policy Library rebuilt with KL-divergence. In the last experiment, we first show the Policy Library rebuilt using KL-divergence in uni-direction and bi-direction of PLKL by executing 50 different tasks in Figure 1 sequentially performed with Library $_{init}$ at starting.

To show that using the KL-divergence as a similarity measure between policies performs better than the way used in PLPR, we review the results of the Rebuilt Library in PLPR firstly. After executing each task, it will append the new policy to the Policy Library if and only if $\max_i W^{-1}_i \infty < W^{-1}_{task} \infty$. Figure 7 shows its result with its known best parameter $= 0.25$[5].

Figure 7: Shadow tasks in Library Rebuilt via PLPR

In Figure 7, the source task distribution generated by PLPR is not uniform: two source policies exist in one room; some rooms have no source policy. The two bad conditions cannot be alleviated together by only adjusting the parameter

Then we show our results of the Policy Library rebuilt by using PLKL. The PLKL method is good for screening the unique policy of each room in both uni-direction way and bi-direction way, as shown in Figure 8,

In Figure 8, to compare the Policy Library generated by the bi-direction KL-divergence (bi-KL) with the uni-KL: some source policies have been replaced by new policies in a way. If the threshold is higher, the source Policy Library after the reconfiguration will be incomplete in some conditions, and if the threshold is lower, the source policy will be replaced many times, of course, this does not affect the availability of the Policy Library.

In this experiment, we also show the comparison of learning effect among different Policy Libraries used for solving a new task. We choose task $_{47}$ to learn. And we set different conditions of initial policy library:

(1) a normal Policy Library $_{init}$ (also the Library with uni-KL)