

Toward Robust Long Range Policy Transfer

Wei-Cheng Tseng¹, Jin-Siang Lin¹, Yao-Min Feng¹, Min Sun^{1,2,3}

¹National Tsing Hua University

²Appier Inc., Taiwan

³MOST Joint Research Center for AI Technology and All Vista Healthcare, Taiwan
{weichengtseng, linjinsiang, yaominlouis}@gapp.nthu.edu.tw, sunmin@ee.nthu.edu.tw

Abstract

Humans can master a new task within a few trials by drawing upon skills acquired through prior experience. To mimic this capability, hierarchical models combining primitive policies learned from prior tasks have been proposed. However, these methods fall short comparing to the human’s range of transferability. We propose a method, which leverages the hierarchical structure to train the combination function and adapt the set of diverse primitive policies alternatively, to efficiently produce a range of complex behaviors on challenging new tasks. We also design two regularization terms to improve the diversity and utilization rate of the primitives in the pre-training phase. We demonstrate that our method outperforms other recent policy transfer methods by combining and adapting these reusable primitives in tasks with continuous action space. The experiment results further show that our approach provides a broader transferring range. The ablation study also show the regularization terms are critical for long range policy transfer. Finally, we show that our method consistently outperforms other methods when the quality of the primitives varies.

1 Introduction

Reinforcement learning (RL) has lots of success in various applications, such as game playing (Brockman et al. 2016; Silver et al. 2017; Mnih et al. 2015), robotics control (Tassa et al. 2018), molecule design (You et al. 2018), and computer system optimization (Mao et al. 2019a,b). Typically, researchers use RL to solve each task independently and from scratch, which makes RL confronted with low sample efficiency. However, compared with humans, the transferability of RL is limited. Especially, humans can learn to solve complex continuous problems (both state space and action space are continuous) efficiently by utilizing prior knowledge. In this work, we want agents to efficiently solve the complex continuous problem by exploiting prior experiences that provide structured exploration based on effective representation.

To this end, we formulate transfer learning in RL as following. We train a policy with one of the RL optimization strategies on the pre-training task. Then, we intend to leverage the policy to master the transferring task. However,

transfer learning in RL may face some fundamental problems. First, unlike supervised learning, the transitions and trajectories are sampled during the training phase based on the interacted policy (Rothfuss et al. 2019). Since the reward distributions are different between the pre-training task and the transferring task, directly finetuning the pre-training policy on transferring tasks may make the agent perform biased structured exploration and get stuck in many low reward trajectories. Second, dynamics shifts between pre-training and transferring tasks may induce the pre-training policy to perform unstructured exploration (Clavera et al. 2019; Nachum et al. 2020). Although domain randomization (Tobin et al. 2017; Nachum et al. 2020) in the pre-training phase may mitigate this problem, we prefer the pre-trained policies to gradually fit the transferring tasks.

Some methods intend to limit dependency between pre-training policy and task-specific information (Goyal et al. 2019; Galashov et al. 2019) by using information bottleneck (Alemi et al. 2017) and variational inference (Kingma and Welling 2013). That way, that pre-training policy does not overfit to a specific task and can be transferred to other tasks. Besides, some methods achieve task transfer by embedding tasks into a latent distribution (Merel et al. 2019; Hausman et al. 2018). However, the latent distribution should be smooth and contain a diverse set of tasks to perform behavior well. Some works propose a hierarchical policy (Frans et al. 2018; Bacon, Harb, and Precup 2017; Peng et al. 2019), which contains a combination function to control how to select or combine a set of primitives. Those works acquire a new selection or combination strategy to control the primitives to master the transferring task if we attain a set of task-agnostic primitives. We find that hierarchical architecture has the potential to enable a better transferring range in continuous control problems.

We propose a transfer learning method in RL. Our pre-training method leverages existing hierarchical structure in the policy consisting of a combination function and a set of primitive policies. We also design objectives to encourage the set of primitives to be diverse and more evenly utilized in pre-training tasks. Notice that we do not use reference data since we expect our method to be generally applicable to all tasks. In many cases, such as flying creatures (Won, Park,

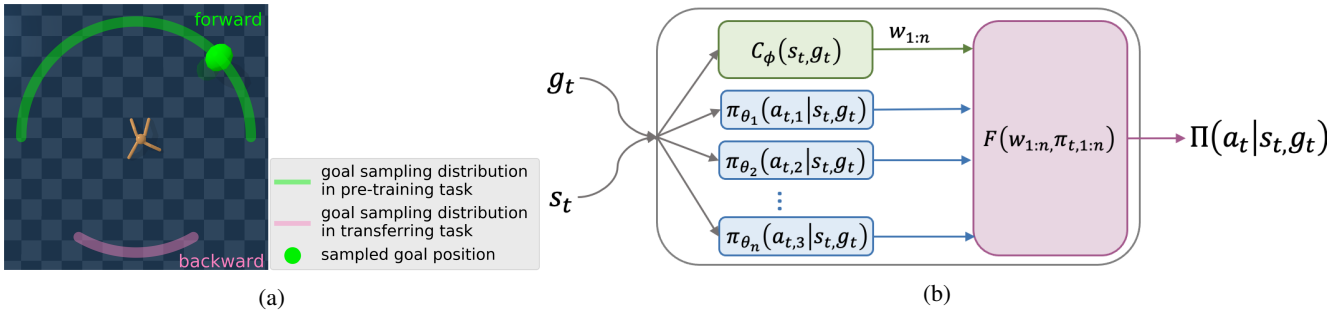


Figure 1: (a) Our motivating example for RL transferring. The green ball represents the target position, which is sampled from the distribution of the task. The goal direction of the pre-training task and transferring task are quite different. (b) The hierarchical policy architecture.

and Lee 2018), Laikago robot¹ or D’Kitty robot², reference data is hard to obtain.

During the transferring phase, we alternatively train the combination function and the primitive policies. This training procedure makes the training not only stable but also flexible in exploration. When training the combination function and freezing primitives in the transferring phase, it utilizes the benefit of the hierarchical structure that abstracts the exploration space. When training the primitives and fixing the combination function, the primitives can be adapted to the transferring task. In our experiments, we demonstrate that training hierarchical policy with our method significantly increases sample efficiency compared to previous work (Peng et al. 2019). Moreover, our method provides a better transferring range. We also provide an ablation study to discuss the effectiveness of our regularization terms. Finally, we show that with different resource constraints on training the pre-training policy, our method still outperforms other methods. The source code is available to the public³.

2 Preliminaries

We consider a multi-task RL framework for transfer learning, consisting of a set of pre-training tasks and transferring tasks. An agent is trained from scratch on the pre-training tasks. Then, it utilizes any skills acquired during pre-training to the transferring tasks. Our objective is to obtain and leverage a set of reusable skills learned from the pre-training tasks to enable the agent to efficiently explore and be more effective at the following transferring tasks.

We denote s as a state, a as an action, r as a reward, and τ as a trajectory consisting of actions and states. Each task is represented by a dynamics model $s_{t+1} \sim p(s_{t+1}|s_t, a_t)$ and a reward function $r_t = r(s_t, a_t, g)$, where g is the task-specific goal such as the target location that an agent intends to reach and a terrain that an agent needs to pass. In multi-task RL, goals $\{g\}$ are sampled from a distribution $p(g)$.

Given a goal g , a trajectory $\tau = \{s_0, a_0, s_1, \dots, s_T\}$ with time horizon T is sampled from a policy $\pi(a|s, g)$. Our objective is to learn an optimal policy π^* that maximizes its

expected return $J(\pi) = \mathbb{E}_{g \sim p(g), \tau \sim p_\pi(\tau|g)}[\sum_{t=0}^T \gamma^t r_t]$ over the distribution of goals $p(g)$ and trajectories $p_\pi(\tau|g)$, where $\gamma \in [0, 1]$ is the discount factor. The probability of the trajectory τ is calculated as follow

$$p_\pi(\tau|g) = p(s_0) \prod_{t=0}^{T-1} p(s_{t+1}|s_t, a_t) \pi(a_t|s_t, g) \quad (1)$$

where $p(s_0)$ is the probability of the initial state s_0 . In transferring learning, despite the same state and action space, the goal distributions, reward functions, and dynamics models in pre-training and transferring tasks are subjected to be different. The difference between the pre-training and transferring tasks is referred to as the range of transfer. Note that a successful transfer can’t be expected for totally unrelated tasks. We consider the scenario where the pre-training tasks can make the agent to learn relevant information of the following transferring tasks, but may not cover the entire set of skills which are useful at the transferring tasks.

3 Method

We will first describe a motivating example in Sec. 3.1. Then we introduce our method in Sec. 3.2. Finally, we show how to apply our method to the existing hierarchical policy in Sec. 3.3.

3.1 A Motivating Example

Let’s consider a model-free RL scenario. We have a pre-training task that an ant needs to achieve a goal position to get some reward, and the goal position is sampled from a half-circle in front of the ant. However, we change the sample distribution of goal positions in the transferring task to a small arc in the back of the ant that does not overlap with the goal positions of the pre-training task (see Fig. 1a). Intuitively speaking, this is a challenging transferring scenario.

There are two straightforward methods to tackle this problem. One is directly finetuning the pre-training policy on the transferring task. However, the pre-trained policy is affected by the goal distribution in the pre-training phase to move forward, and this conflicts with the goal distribution in transferring task to move backward. Therefore, directly finetuning may corrupt the information learned from the pre-training task, which is known as the catastrophic forgetting.

¹<http://www.unitree.cc/e/action/ShowInfo.php?classid=6&id=1>

²<https://www.trossenrobotics.com/d-kitty.aspx>

³https://weichengtseng.github.io/project_website/aaai21

The other is to train a new policy from scratch on the transferring task. Training from scratch can saturate to a good point, but it may need lots of trails. Our proposed method aims to address the drawbacks of these two basic methods.

Algorithm 1: Full Algorithm

```

// pre-training phase ;
initialize  $\theta_{1:n}$  and  $\phi$  ;
let  $J_{RL}(\theta_{1:k}, \phi)$  be the objective function of some
specific RL optimization ;
while not converge do
  train combination function  $\phi$  and primitives  $\theta_{1:n}$ 
  with  $J_{pre}(\theta_{1:k}, \phi) =$ 
   $J_{RL}(\theta_{1:k}, \phi) + \alpha \cdot J(\pi_{1:k}) + \beta \cdot J(w_{1:k})$ 
// transferring phase ;
reinitialize  $\phi$  ;
while not converge do
  Disable the gradient of primitives ;
  Enable the gradient of combination function ;
  for  $i = 1 : p$  do
    train combination function  $\phi$  with
     $J_{transfer}(\theta_{1:k}, \phi) = J_{RL}(\theta_{1:k}, \phi)$ 
  Disable the gradient of combination function ;
  Enable the gradient of primitives ;
  for  $i = 1 : p$  do
    train primitives  $\theta_{1:n}$  with
     $J_{transfer}(\theta_{1:k}, \phi) = J_{RL}(\theta_{1:k}, \phi)$ 

```

3.2 Our Method

We describe our method in this section. Our policy architecture contains: 1) A set of primitive policies $\pi_{\theta_1}(a|s, g), \pi_{\theta_2}(a|s, g), \dots, \pi_{\theta_n}(a|s, g)$ with parameters $\theta_{1:n}$, and each primitive is an independent policy that outputs action distribution conditioned on s and g . 2) A combination function $C_\phi(s, g)$ with parameter ϕ outputs weight $w_{i:n}$, where w_i specifies the importance of primitive π_{θ_i} . $F(\pi_{1:n}, w_{1:n})$ specifies how to combine primitives with specified weight $w_{1:n}$. Typically, the larger the weight w_i , the more contributions from primitive π_{θ_i} .

During the pre-training phase, the hierarchical policy is end-to-end trained with an off-the-shelf RL optimization method. Our goal is to learn a set of primitives such that the combination function can compose them to form a complete behavior. However, not all primitives learned in pre-training are equally good for long range transfer. We identify two issues to be addressed. First, diversity of primitives is essential to improve the transferability. Without encouraging diversity, some of the primitives may learn similar behavior which better fits the pre-training tasks but affects the transferability. Second, when more primitives are introduced to compose more complex behaviors and improve the transferability, the utilization rate of each primitive varies more. In some extreme cases, some primitives may seldom or never used in the pre-training phase, so these primitives are not updated by RL optimization. Hence, it is important to en-

courage the utilization rates of primitives to be more evenly distributed.

To mitigate these two issues, we propose two regularization terms. We introduce the first regularization term that separates the distributions of the primitives from each other's so that the primitives become diverse. The intuitive idea to measure the difference between two probability distributions is calculating KL divergence. Therefore, we calculate the average of arbitrary pair of primitives as our regularization term, and we call it Diversity Regularization (DR)

$$\begin{aligned}
J(\pi_{\theta_{1:k}}) &= \frac{1}{k(k-1)} \sum_{i \neq j} D_{KL}(\pi_{\theta_i} | \pi_{\theta_j}) \\
&= \frac{2}{k(k-1)} \sum_{i \neq j, i < j} D_{JS}(\pi_{\theta_i} | \pi_{\theta_j})
\end{aligned} \tag{2}$$

Then, we propose another regularization term to encourage utilization rates of primitives to be more evenly used in the pre-training task. We try to model the weight of primitives as categorical distribution and use the entropy of distribution as our regularization term, and we call it Utility Regularization (UR).

$$J(w_{1:k}) = - \sum_i \left(\frac{e^{w_i}}{\sum_j e^{w_j}} \right) \log \frac{e^{w_i}}{\sum_j e^{w_j}} \tag{3}$$

Therefore, the overall objective in the pre-training phase is shown as below.

$$J_{pre}(\theta_{1:k}, \phi) = J_{RL}(\theta_{1:k}, \phi) + \alpha \cdot J(\pi_{1:k}) + \beta \cdot J(w_{1:k}) \tag{4}$$

where α and β are the hyperparameters and $J_{RL}(\theta_{1:k}, \phi)$ can be any reinforcement learning optimization method.

During the transferring phase, we reinitialize the weights of the combination function and leverage the primitive directly from pre-trained in the pre-training task. We first update the combination function and freeze the primitives. This views the combination function as the policy that learns to combine the primitives to master the transferring task. As illustrated in the motivating example, the range of transfer between the pre-training and transferring tasks is likely to limit the performance of only training the combination function. Therefore, after p iterations, we switch to finetune the primitives with the combination function fixed. This method makes the primitives become more applicable to the transferring task. After p iterations, we freeze primitives and train combination function again. This is to prevent the skills in the primitives to be severely forgotten during finetuning. The strategy is repeated several times until the hierarchical policy converges (see algorithm 1).

3.3 Applying Our Method to the Policy Architecture

We leverage the multiplicative combination rule (Peng et al. 2019)

$$F(\pi_{\theta_{1:n}}, w_{1:n}) = \frac{1}{Z(s, g)} \prod_{i=1}^k \pi_{\theta_i}(a|s)^{w_i} \tag{5}$$

where $\pi_{\theta_{1:n}}(a|s)$ is the primitive policies and $w_{1:n}$ is generated from combination function $C_\phi(s, g)$. $Z(s, g)$ is the

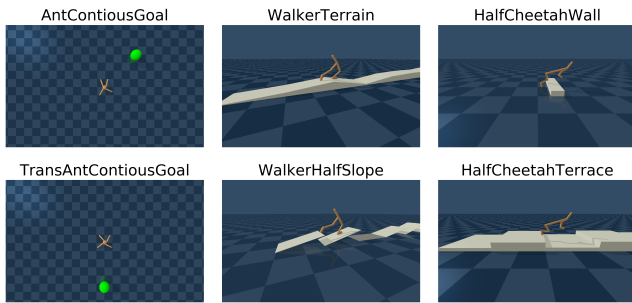


Figure 2: Environments used to evaluate our method. The first row is the pre-training tasks, and the second row is the transferring tasks. All the baseline and our method are evaluated with the three agents: ant, 2dwalker, halfcheetah.

partition function that ensures the composite distribution is normalized. $F(\pi_{\theta_{1:n}}, w_{1:n})$ multiplies the primitive policies along with their corresponding weights. The weights determine the importance of each primitive policies to compose the action distribution at a time step, with a larger weight representing a larger influence. Note that to make the primitives task-agnostic (i.e., more transferable), we restrict the primitives $\pi_{\theta_{1:n}}(a|s)$ to only get s . During the pre-training phase, the combination function and primitive policies are trained in an end-to-end manner. During the transferring phase, we train the combination function and primitive policies alternatively as described in the previous section.

4 Related Works

Learning meaningful and reusable representations that can be transferred across multiple tasks is a popular research direction in machine learning (Argyriou, Evgeniou, and Pontil 2007). One of the straightforward transfer learning methods is finetuning. That is, a network is first trained on a source domain. Then, the learned representation or features are reused in another domain by finetuning via backpropagation (Hinton and Salakhutdinov 2006; Kemker et al. 2018; Kirkpatrick et al. 2016). However, backpropagation may destroy previously learned representations or features before the network leverages them in the target domain.

4.1 Information-Based Methods

Some methods try to learn a policy that doesn't highly depend on task-specific information during pre-training and only acquire task-specific information when the policy needs to make a critical decision to solve the task (Goyal et al. 2019; Galashov et al. 2019; Goyal et al. 2020). These methods may leverage some concepts like information bottleneck (Alemi et al. 2017) and variational inference (Kingma and Welling 2013) to restrict policy and task-specific information dependency. Since the policy only gets limited task-specific information, the policy can potentially be transferred to other tasks.

4.2 Hierarchical Methods

Some methods intend to use a hierarchical policy that typically contains a master policy or combination function and

a set of low-level policies (Bacon, Harb, and Precup 2017). Then, by training the hierarchical policy in the pre-training task and acquire a set of primitives, we can reschedule or re-compose these to master transferring tasks. The combination function in some method only activates one sub-policy at a time (Liu and Hodgins 2017; Peng et al. 2018; Frans et al. 2018) while the other combination function allows multiple primitive to be executed at the same time (Peng et al. 2019; Qureshi et al. 2020). Another kind of methods leverages some ideas from hierarchical reinforcement learning (HRL) (Nachum et al. 2018; Levy, Platt, and Saenko 2019). The policy architecture typically contain a high-level policy and a low-level policy. The high-level policy outputs a goal state or condition to the low-level policy, and the low-level policy intends to achieve the given goal state. Some works (Li et al. 2019, 2020) propose to finetune policy with HRL framework to perform transfer learning in RL.

4.3 Latent Space Methods

Some methods specify actions through the latent representation which is further transformed to the actions of the underlying system (Merel et al. 2019; Burgard, Brock, and Stachniss 2008; Chandak et al. 2019). Therefore, a common approach is getting the latent space by pre-training the policy on the pre-training task, and further transfer the policy to the downstream task (Haarnoja et al. 2018). To encourage the latent space is diverse enough, one solution is to leverage reference data⁴ to pre-train the latent space (Merel et al. 2019). Other diversity-driven pre-training methods have been proposed to help the latent space to form semantically various behaviors (Hausman et al. 2018).

5 Experiments

In this section, we introduce the evaluation tasks in Sec. 5.1 and list the baselines that we intend to compare in Sec. 5.2. The results of the methods evaluated in our environments are discussed in Sec. 5.3. Aside from option-critic (Bacon, Harb, and Precup 2017), all the experiments are trained with PPO (Schulman et al. 2017) and Generalized Advantage Estimation (GAE) (Schulman et al. 2016). The detailed hyperparameter settings are shown in supplementary Sec. 4.1. We further show that our method has a broader transferring range compared with other baselines in Sec. 5.4. Then, we demonstrate the effectiveness of the regularization terms in Sec. 5.5. Finally, we demonstrate that our method can perform well even if the primitive policies are not good enough in Sec. 5.6. Some addition proprieties of our method are further discussed in supplementary Sec. 2.

5.1 Tasks

We consider three agents (see Fig. 2): quadruped (ant) with 12 DoF and 8 actuators; 2dwalker with 6 DoF and 6 actuators; halfcheetah with 6 DoF and 6 actuators. All the tasks are built with PyBullet⁵. The implementation detail of these tasks will be described in supplementary Sec. 4.2. and we discuss the tasks in the following sections.

⁴<http://mocap.cs.sfu.ca> and <http://mocap.cs.cmu.edu>

⁵<http://pybullet.org>




Figure 3: Performance of different transferring methods. For better visualization, we use exponential moving average to smooth the learning curve, and each learning curve is grouped with ten random seeds. The transparent part represents the maximum and minimum of the learning curves. From these figures, we show that our method achieves better performance than other methods.

Pre-Training Tasks

- **AntContinuousGoal:** An ant needs to move to the target position which is the task-specific information g , and the target direction is sampled from $[-0.5\pi, 0.5\pi]$ with radius 5 as we showed in Fig. 1a. Once the ant reaches the goal position, the goal position will be re-sampled in the same way.
- **WalkerTerrain:** A 2d-walker needs to move forward on terrain, and the slope of the terrain is sampled from a specific range. The task-specific information g is the terrain in front of the agent. Therefore, the 2d walker needs to learn how to walk smoothly on planes with different slopes.
- **HalfCheetahWall:** A halfcheetah needs to move forward, and there is a wall every 3 to 5 meters. The height of the wall is sampled from a specific range. The task-specific information g is the terrain in front of the agent. Therefore, the halfcheetah needs to climb or jump over the wall.

Transferring Tasks

- **TransAntContinuousGoal:** An ant needs to move to the target position, and the target position is sampled from $[5\pi/6, 7\pi/6]$ with radius 5, as we showed in Fig. 1a. Note that the target direction doesn't overlap with that of the pre-training task. Once the ant reaches the goal position, the goal position will be re-sampled in the same way. Therefore, in this case, the difference between pre-training task and transferring task is goal distribution $p(g)$.
- **WalkerHalfSlope:** A 2d walker needs to move forward in terrain, but there are cliffs between each plane. Therefore, the 2d walker should be robust to these cliffs. Therefore, in this case, the differences are goal distribution $p(g)$ and dynamics $p(s_{t+1}|s_t, a_t)$.
- **HalfCheetahTerrace:** A halfcheetah needs to move forward on a terrace that is formed with lots of horizontal platforms with target speed, so it needs to jump or climb up to the higher platform and does not fall to the lower platform. The difference in height between two platforms is sampled from a specific range. Therefore, in this case,

the differences are goal distribution $p(g)$ and dynamics $p(s_{t+1}|s_t, a_t)$.

5.2 Baselines

We define the baselines that will be discussed in the following sections, and the implementation detail will be described in supplementation material Sec. 4.1.

Scratch: We directly train a policy with PPO(Schulman et al. 2017) on the transferring task, and it is one of the most straightforward methods to tackle a task. **Finetune:** We first train a policy in the pre-training task. Then, we directly finetune the policy in the transferring task. It is another the most straightforward method to tackle a task. **MCP** (Peng et al. 2019): A multiplicative model that enables the agent to activate multiple primitives simultaneously is trained. Each primitive specializes in different behaviors that can be combined to span a continuous spectrum of skills on the pre-training task. During the transferring phase, the primitives are fixed, and only the combination function is updated. **Option-Critic** (Bacon, Harb, and Precup 2017): Both intra-option policies and termination conditions of options are learned, followed by the policy over options, and without specifying any additional rewards or subgoals. The options are assigned to perform sequentially. One option works for several timesteps until being stopped by the termination function. The entire network is pre-trained on the pre-training task. Then, it is directly finetuned to the transferring task. **MLSH** (Frans et al. 2018): A hierarchical policy where the master policy switches between a set of sub-policies is learned. The master policy chooses a sub-policy every N timesteps, and the selected sub-policy is then executed by the agent for N timesteps to interact with the environment to constitute a high-level action. Both master policy and sub-policies are optimized on the pre-training task. During the transferring phase, we freeze the sub-policies and only train the master policy, which makes the master policy learn how to utilize the fixed primitives.

Aside from the baselines mentioned above, some additional baselines are discussed in supplementary Sec. 1.

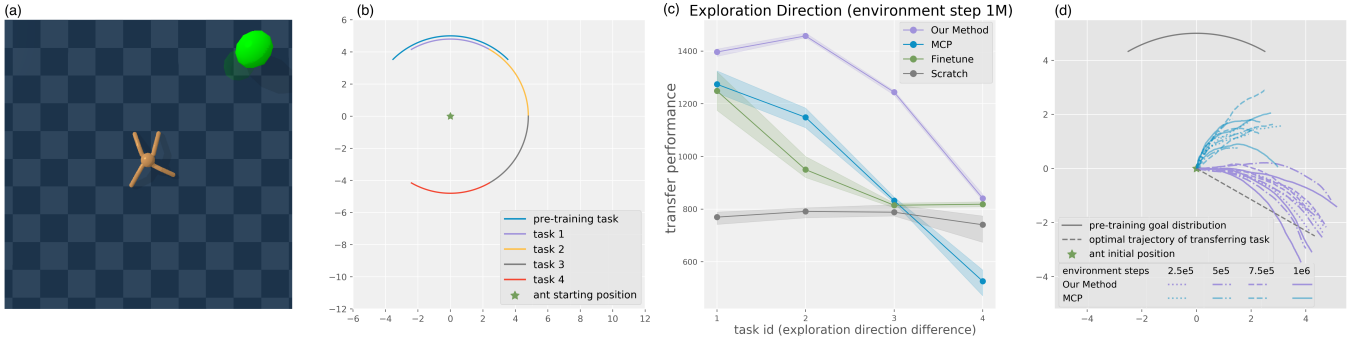


Figure 4: Performance of different transferring methods. From these figures, we show that our method achieves better performance than other methods. (a): the experiment is based on the ant agent. (b): the goal position sampling range for each task. From task 1 to task 4, the task difference between transferring task and pre-training task becomes larger. (c): the result of each method on different transferring tasks. (d): the trajectories of our method and MCP under different numbers of training iterations. It is clear that our method was not affected by the pre-training task compared with MCP.




Figure 5: Normalized performance with different period p . The performance is normalized by the best average return achieved in each task so that we can plot the result of all the tasks in a single plot. We find that the performance is not affected much by the period p .

5.3 Comparisons with Baseline

We run our method and other baselines in the three continuous problems described in the previous section. The network architecture, hyperparameter setting, and other implementation detail will be described in supplementary Sec. 4.1. For Fig. 3, we find that our method outperforms other methods in the transferring phase. MCP tends to perform well at the beginning of training, and this phenomenon may be caused by hierarchical abstraction. However, freezing primitive policies may restrict the transferring range of MCP, which induces that MCP cannot converge to a better result. Since our method acquires a set of distinct and even utilization rate primitives and allows the primitives to adapt to the transferring task, it achieves significantly better reward the fastest among all the baseline algorithms. With the knowledge learned from the pretraining tasks, finetuning performs better than training from scratch. This is because training from scratch is required to learn everything, including task distribution and dynamics. However, our method significantly outperforms finetuning by striking a better trade-off between combining the primitives to efficiently exploring

the new task and gradually adapting the primitive to the new task. MLSH and option-critic do not perform well in our transfer task since they choose the primitives serially, which makes the primitives not decomposed well. When the task distribution and dynamics change in the transferring phase, they are no longer able to provide suitable behavior.

One concern about our method is that it introduces a hyperparameter p , which controls the period of the number of the gradient steps before switching between primitives and combination function. To this end, we show the performance of our method with different period p in Fig. 5. We find that our method is not sensitive to period p , so there is not much effort to tune this hyperparameter.

5.4 Transferring Range

The transferring range is critical in transfer learning. To demonstrate the transferring range of each method, we redesign the goal position of AntContinuousGoalEnv (see Fig. 4 (a, b)). The goal position pre-training task is sampled from an arc where the center angle is $[-\pi/4, \pi/4]$ and radius 5 meters. As for transferring tasks, we design four transferring tasks, and the goal position of these four transferring tasks are sampled from $[-\pi/6, \pi/6]$, $[\pi/6, 3\pi/6]$, $[3\pi/6, 5\pi/6]$ and $[5\pi/6, 7\pi/6]$, respectively. In other words, the four transferring tasks are ordered by the scale of the task difference between the pre-training task and the transferring task.

In Fig. 4 (c), we find that our method has a larger transferring range compared with other methods. Note that we report the performance at 1 million environment steps. All the other transferring methods get worse as exploration direction becomes different. As the exploration direction difference increases, MCP may perform worse than finetuning. It may be caused by fixing the primitive policies, which may limit the ability to adapt to the transferring task. Besides, if we plot the trajectory of the ant in the transferring task 3 (see Fig. 4 (d)) with a fixed goal position, we find MCP is biased by the pre-training task. Though the ant mitigates this bias by more training iteration, it tends to move forward and turn to the goal direction. In contrast, since our method al-




Figure 6: We evaluate the effectiveness of the pre-trained policy by using different checkpoints with different performance. We find that our method outperforms other methods given different pre-training performances.

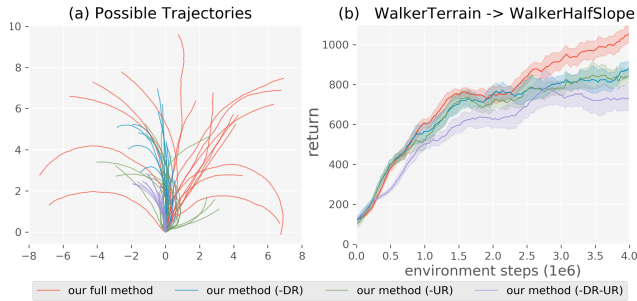


Figure 7: Ablation study on our method. (a) Trajectories that generate with random sample $w_{1:k}$. (b) We evaluate the our method with and without the regularization terms in the walker scenario.

lows the primitive policy to adapt to transferring task, the ant tends to act in goal direction directly. The experiment result implies that our method has a better transferring range than other methods.

5.5 Ablation Study

To evaluate whether the regularization terms (DR and UR) are effective, we conduct the ablation study on the ant and walker scenario. For the ant scenario, we pre-train the primitives with and without the regularization terms. Then, we plot the trajectories with random sampled combination weights $w_{1:k}$ in Fig. 7. We can find that primitives trained with both DR and UR help the agent move in a different direction with longer distance. To be more specific, the trajectories with only DR tend to be short but diverse in the moving direction. It is because DR makes the primitive difference between each other. However, since DR does not enhance the utilization rate of primitives, some primitives become redundant and affect the entire policy’s performance (i.e., shorter trajectories). On the other hand, the trajectories with only UR tend to be long but lack of diversity in the moving direction. Note that the long trajectories are generated even with random weights in the combination function. This is because, with UR, RL optimization more evenly updates all the primitives during pre-training. With both DR and UR, long and diverse trajectories enables structure exploration which helps the combination function to be effi-

ciently retrained during the transferring phase.

We further conduct pre-training and transferring experiments in walker scenario. From Fig. 7 (b), we can find that both diversity regularization and utility regularization can the policy achieve better performance with better sample efficiency. Our method with all regularization terms outperform other versions. More ablation study are provide in supplementary Sec. 1.

5.6 Quality of the Primitives

We further study how the quality of the primitives on the pre-training task will affect policy transfer performance. A stable policy transfer method should still perform reasonably well even when the quality of the primitives varies. To be more specific, we train the primitives with different training iteration in pre-training. From Fig. 6, we observe that our method achieves better performance in the transferring phase compared to other baselines. The performance of MCP is significantly affected by the quality of the primitives as the primitives are fixed. We also observed that, in the second and third tasks, the better the primitives, the better the our method performance in the transferring task. However, in the first task where the goal distributions are the opposite between the pre-training and transferring tasks, the better the primitives do not imply the better the our method performance in the transferring task. Although we found that we prefer the different quality of the primitives depending on the range of transfer, our method is the most stable method outperforming all baseline methods consistently.

6 Conclusion

We propose a method that leverages hierarchical structures by training over different function combinations with two novel regularization terms and adapting primitive policies alternatively. The experiments show that our approach outperforms previous policy transfer methods under the long range transferring scenario for continuous action spaces. We also empirically show that our method provides a larger transferring range as well as an effective adaptation by varying the scale of the task difference between the pre-training task and transferring task. The ablation study also provides evidence of the effectiveness of our regularization terms. Finally, our method achieves more stable performance than other methods do when the quality of the primitive varies.

Acknowledgments

This project is funded by Ministry of Science and Technology of Taiwan (MOST 109-2634-F-007-016, MOST Joint Research Center for AI Technology and All Vista Healthcare).

References

- Alemi, A. A.; Fischer, I. C.; Dillon, J. V.; and Murphy, K. 2017. Deep Variational Information Bottleneck. In *International Conference on Learning Representations*.
- Argyriou, A.; Evgeniou, T.; and Pontil, M. 2007. Multi-Task Feature Learning. In *Advances in Neural Information Processing Systems 19*.
- Bacon, P.-L.; Harb, J.; and Precup, D. 2017. The Option-Critic Architecture. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*.
- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. OpenAI Gym.
- Burgard, W.; Brock, O.; and Stachniss, C. 2008. *Learning Omnidirectional Path Following Using Dimensionality Reduction*.
- Chandak, Y.; Theocharous, G.; Kostas, J.; Jordan, S.; and Thomas, P. 2019. Learning Action Representations for Reinforcement Learning. In *Proceedings of the 36th International Conference on Machine Learning*.
- Clavera, I.; Nagabandi, A.; Liu, S.; Fearing, R. S.; Abbeel, P.; Levine, S.; and Finn, C. 2019. Learning to Adapt in Dynamic, Real-World Environments through Meta-Reinforcement Learning. In *International Conference on Learning Representations*.
- Frans, K.; Ho, J.; Chen, X.; Abbeel, P.; and Schulman, J. 2018. META LEARNING SHARED HIERARCHIES. In *International Conference on Learning Representations*.
- Galashov, A.; Jayakumar, S.; Hasenclever, L.; Tirumala, D.; Schwarz, J.; Desjardins, G.; Czarnecki, W. M.; Teh, Y. W.; Pascanu, R.; and Heess, N. 2019. Information asymmetry in KL-regularized RL. In *International Conference on Learning Representations*.
- Goyal, A.; Islam, R.; Strouse, D.; Ahmed, Z.; Larochelle, H.; Botvinick, M.; Levine, S.; and Bengio, Y. 2019. Transfer and Exploration via the Information Bottleneck. In *International Conference on Learning Representations*.
- Goyal, A.; Sodhani, S.; Binas, J.; Peng, X. B.; Levine, S.; and Bengio, Y. 2020. Reinforcement Learning with Competitive Ensembles of Information-Constrained Primitives. In *International Conference on Learning Representations*.
- Haarnoja, T.; Hartikainen, K.; Abbeel, P.; and Levine, S. 2018. Latent Space Policies for Hierarchical Reinforcement Learning. In *Proceedings of the 35th International Conference on Machine Learning*.
- Hausman, K.; Springenberg, J. T.; Wang, Z.; Heess, N.; and Riedmiller, M. 2018. Learning an Embedding Space for Transferable Robot Skills. In *International Conference on Learning Representations*.
- Hinton, G. E.; and Salakhutdinov, R. R. 2006. Reducing the Dimensionality of Data with Neural Networks. *Science* 313(5786): 504–507. ISSN 0036-8075. doi: 10.1126/science.1127647. URL <https://science.sciencemag.org/content/313/5786/504>.
- Kemker, R.; Abitino, A.; McClure, M.; and Kanan, C. 2018. Measuring Catastrophic Forgetting in Neural Networks. In *AAAI*.
- Kingma, D. P.; and Welling, M. 2013. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations*.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N. C.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; Hassabis, D.; Clopath, C.; Kumaran, D.; and Hadsell, R. 2016. Overcoming catastrophic forgetting in neural networks. *CoRR* abs/1612.00796. URL <http://arxiv.org/abs/1612.00796>.
- Levy, A.; Platt, R.; and Saenko, K. 2019. Hierarchical Reinforcement Learning with Hindsight. In *International Conference on Learning Representations*.
- Li, A.; Florensa, C.; Clavera, I.; and Abbeel, P. 2020. Sub-policy Adaptation for Hierarchical Reinforcement Learning. In *International Conference on Learning Representations*.
- Li, S.; Wang, R.; Tang, M.; and Zhang, C. 2019. Hierarchical Reinforcement Learning with Advantage-Based Auxiliary Rewards. In *Advances in Neural Information Processing Systems 32*.
- Liu, L.; and Hodgins, J. 2017. Learning to Schedule Control Fragments for Physics-Based Characters Using Deep Q-Learning. *ACM Trans. Graph.* 36(4). ISSN 0730-0301.
- Mao, H.; Negi, P.; Narayan, A.; Wang, H.; Yang, J.; Wang, H.; Marcus, R.; addanki, r.; Khani Shirkoohi, M.; He, S.; Nathan, V.; Cangialosi, F.; Venkatakrishnan, S.; Weng, W.-H.; Han, S.; Kraska, T.; and Alizadeh, D. 2019a. Park: An Open Platform for Learning-Augmented Computer Systems. In *Advances in Neural Information Processing Systems 32*.
- Mao, H.; Venkatakrishnan, S. B.; Schwarzkopf, M.; and Alizadeh, M. 2019b. Variance Reduction for Reinforcement Learning in Input-Driven Environments. In *International Conference on Learning Representations*.
- Merel, J.; Hasenclever, L.; Galashov, A.; Ahuja, A.; Pham, V.; Wayne, G.; Teh, Y. W.; and Heess, N. 2019. Neural Probabilistic Motor Primitives for Humanoid Control. In *International Conference on Learning Representations*.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A.; Veness, J.; Bellemare, M.; Graves, A.; Riedmiller, M.; Fidjeland, A.; Ostrovski, G.; Petersen, S.; Beattie, C.; Sadik, A.; Antonoglou, I.; King, H.; Kumaran, D.; Wierstra, D.; Legg, S.; and Hassabis, D. 2015. Human-level control through deep reinforcement learning. *Nature* 518.
- Nachum, O.; Ahn, M.; Ponte, H.; Gu, S. S.; and Kumar, V. 2020. Multi-Agent Manipulation via Locomotion using Hierarchical Sim2Real. In *Proceedings of the Conference on Robot Learning*.

Nachum, O.; Gu, S. S.; Lee, H.; and Levine, S. 2018. Data-Efficient Hierarchical Reinforcement Learning. In *Advances in Neural Information Processing Systems 31*.

Peng, X. B.; Abbeel, P.; Levine, S.; and van de Panne, M. 2018. DeepMimic: Example-guided Deep Reinforcement Learning of Physics-based Character Skills. *ACM Trans. Graph.* 37(4): 143:1–143:14. ISSN 0730-0301.

Peng, X. B.; Chang, M.; Zhang, G.; Abbeel, P.; and Levine, S. 2019. MCP: Learning Composable Hierarchical Control with Multiplicative Compositional Policies. In *Advances in Neural Information Processing Systems 32*.

Qureshi, A. H.; Johnson, J. J.; Qin, Y.; Henderson, T.; Boots, B.; and Yip, M. C. 2020. Composing Task-Agnostic Policies with Deep Reinforcement Learning. In *International Conference on Learning Representations*.

Rothfuss, J.; Lee, D.; Clavera, I.; Asfour, T.; and Abbeel, P. 2019. ProMP: Proximal Meta-Policy Search. In *International Conference on Learning Representations*.

Schulman, J.; Moritz, P.; Levine, S.; Jordan, M.; and Abbeel, P. 2016. High-Dimensional Continuous Control Using Generalized Advantage Estimation. In *Proceedings of the International Conference on Learning Representations (ICLR)*.

Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. *CoRR* abs/1707.06347. URL <http://arxiv.org/abs/1707.06347>.

Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; Bolton, A.; Chen, Y.; Lillicrap, T.; Hui, F.; Sifre, L.; Driessche, G.; Graepel, T.; and Hassabis, D. 2017. Mastering the game of Go without human knowledge. *Nature* 550.

Tassa, Y.; Doron, Y.; Muldal, A.; Erez, T.; Li, Y.; de Las Casas, D.; Budden, D.; Abdolmaleki, A.; Merel, J.; Lefrancq, A.; Lillicrap, T.; and Riedmiller, M. 2018. DeepMind Control Suite. Technical report, DeepMind.

Tobin, J.; Fong, R.; Ray, A.; Schneider, J.; Zaremba, W.; and Abbeel, P. 2017. Domain randomization for transferring deep neural networks from simulation to the real world. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.

Won, J.; Park, J.; and Lee, J. 2018. Aerobatics Control of Flying Creatures via Self-Regulated Learning. *ACM Trans. Graph.* 37(6). ISSN 0730-0301. doi:10.1145/3272127.3275023. URL <https://doi.org/10.1145/3272127.3275023>.

You, J.; Liu, B.; Ying, Z.; Pande, V.; and Leskovec, J. 2018. Graph Convolutional Policy Network for Goal-Directed Molecular Graph Generation. In *Advances in Neural Information Processing Systems 31*.