

Received May 19, 2021, accepted June 5, 2021, date of publication June 18, 2021, date of current version July 15, 2021.

Digital Object Identifier 10.1109/ACCESS.2021.3090364

# Subgoal-Based Reward Shaping to Improve Efficiency in Reinforcement Learning

TAKATO OKUDO <sup>ORCID</sup> AND SEIJI YAMADA <sup>ORCID</sup>, (Member, IEEE)

Department of Informatics, The Graduate University for Advanced Studies, SOKENDAI, Tokyo 101-8430, Japan  
National Institute of Informatics, Tokyo 101-8430, Japan

Corresponding author: Takato Okudo (okudo@nii.ac.jp)

**ABSTRACT** Reinforcement learning, which acquires a policy maximizing long-term rewards, has been actively studied. Unfortunately, this learning type is too slow and difficult to use in practical situations because the state-action space becomes huge in real environments. Many studies have incorporated human knowledge into reinforcement Learning. Though human knowledge on trajectories is often used, a human could be asked to control an AI agent, which can be difficult. Knowledge on subgoals may lessen this requirement because humans need only to consider a few representative states on an optimal trajectory in their minds. The essential factor for learning efficiency is rewards. Potential-based reward shaping is a basic method for enriching rewards. However, it is often difficult to incorporate subgoals for accelerating learning over potential-based reward shaping. This is because the appropriate potentials are not intuitive for humans. We extend potential-based reward shaping and propose a subgoal-based reward shaping. The method makes it easier for human trainers to share their knowledge of subgoals. To evaluate our method, we obtained a subgoal series from participants and conducted experiments in three domains, four-rooms(discrete states and discrete actions), pinball(continuous and discrete), and picking(both continuous). We compared our method with a baseline reinforcement learning algorithm and other subgoal-based methods, including random subgoal and naive subgoal-based reward shaping. As a result, we found out that our reward shaping outperformed all other methods in learning efficiency.

**INDEX TERMS** Reinforcement learning, subgoals as human knowledge, potential-based reward shaping, reward shaping.

## I. INTRODUCTION

Reinforcement learning(RL) can acquire a policy maximizing long-term rewards in an environment. Designers do not need to specify how to achieve a goal; they only need to specify what a learning agent should achieve with a reward function. A reinforcement learning agent performs both exploration and exploitation to find how to achieve a goal by itself. It is common for the state-action space to be quite large in a real environment like robotics. As the space becomes larger, the number of iterations to learn the optimal policies exponentially increases, and the learning becomes too slow to obtain the optimal policies in a realistic amount of time. Since a human could have knowledge that would be helpful to such an agent in some cases, a promising approach is utilizing human knowledge [1]. Wang and Taylor's approach transfers

a policy by requesting that a human selects an action for a state during the learning phase [2]. Griffith *et al.* [3] used interactive human feedback as direct policy labels.

The reward function is the most related to learning efficiency. Most difficult tasks in RL have a sparse reward function [4]. The agent is not able to evaluate its policy due to it and to learn the optimal policy. In contrast, learning speeds up when the reward function is dense. Inverse reinforcement learning (IRL) [5] is the most popular method for enriching the reward function. IRL uses an optimal policy to generate a dense reward function. Recent studies have utilized optimal trajectories [6]. There is the question of the cost for the teacher in providing the optimal trajectories or policies. Humans sometimes have difficulty providing these because of the skills they may or may not have. In particular, in robotics tasks, humans are required to have robot-handling skills and knowledge on the optimal trajectory. We focus on the knowledge of subgoals because there are no requirements

The associate editor coordinating the review of this manuscript and approving it for publication was Xiaochun Cheng.

for robot-handling skills. Humans need only to think of a few subgoal states on the optimal trajectory in their minds. Moreover, trajectories that fail can be utilized. Potential-based reward shaping is able to add external rewards while keeping the optimal policy of the environment [7]. It is calculated as the difference between the real-number functions (potential function) of the previous and current state. Generally, a naive idea of the potential function for utilizing the knowledge of subgoals is high potential only for subgoal states. However, the method cannot work well as shown in Section VI, so it is not easy to incorporate the knowledge of subgoals over potential-based reward shaping.

In this paper, we propose subgoal-based reward shaping that accelerates learning by using the knowledge of subgoals with the original optimal policy remaining. Our method focuses on a single agent setting for episodic tasks with MDP with a single constant goal. The strength is that the reward shaping needs only the knowledge of subgoals to accelerate learning. Getting the subgoals does not require handling the agent such as in the case of getting trajectories. Furthermore, trajectories with failure can be utilized by extracting subgoals from them. The main contributions of this work are:

- learning that is accelerating by reward shaping with a sequence of subgoals,
- a user study for acquiring a sequence of subgoals and a comparison between participants' subgoals and ones generated randomly,
- experimental results on three tasks that demonstrate a significant improvement in learning efficiency compared with a baseline algorithm and a traditional method.

The remainder of this paper is organized as the followings. Section II reviews related work on RL and similar studies to our reward shaping with subgoal knowledge. The preliminaries including notations and algorithms for RL and potential based reward shaping are described in Section III. Then, we explain the definition of subgoals, a subgoal-based reward shaping algorithm we proposed in Section IV. The next Section V gives an overview of user study in which we obtained subgoal knowledge from participants. The setting and results of various experiments we conducted for evaluation of our proposed method and descriptions on three domains, four-rooms, pinball, pick and place tasks in Section VI. In Section VII, discussions on subgoals given by participants, hyper parameters tuning and others are discussed. The conclusion is provided in Section VIII.

## II. RELATED WORK

### A. HUMAN KNOWLEDGE FOR REINFORCEMENT LEARNING

There are many studies utilizing human knowledge to RL; trajectory, policy, preference, action, feedback, and subgoal. From policy or trajectory which human provides, IRL infers an unobserved reward function [8]. It is known to be difficult to design the reward functions because of an unspecified reward [9] and a reward hacking [10]. The unspecified

reward is leaving out important aspects in reward design. The reward hacking is short of reward design for penalizing all possible misbehavior. Since the designer only provide the policy or trajectories without defining the reward function directly, the IRL overcomes these difficulties. The reward function is modeled as a linear sum of weighted features, and is acquired by a margin-based optimization [5] or an entropy-based optimization [6]. Since the IRL generates a rich reward function, it is also a powerful solution when an environment has a sparse reward. The approach is same to ours. The difference is that we use only some specific states as subgoal.

In RL with human preferences [11], a human teacher compares a pair of trajectory segments and select the better one than the other one. The method learned a policy without access to the environmental rewards and the human preference predictor.

In dataset aggregation(DAGGER) [12], a human expert selects an action in a state which an agent selects.

In the field of interactive reinforcement learning, a learning agent interacts with a human trainer as well as the environment [13]. The TAMER framework is a typical interactive framework for reinforcement learning [14]. The human trainer observes the agent's actions and provides binary feedback during learning. Since humans often do not have programming skills and knowledge on algorithms, the method relaxes the requirements to be a trainer. We aim for fewer trainer requirements, and we use a GUI on a web system in experiments with navigation tasks.

### B. SUBGOALS FOR REINFORCEMENT LEARNING

We consider the hindsight experience replay (HER) algorithm [15], which uses subgoals. HER regards a failure as a success in hindsight, and it generates a reward for the failure. The algorithm interprets a failure as a process for achieving a target, and it is really similar to a subgoal. HER accelerates learning in a multi-goal environment. We focus on a single-goal environment, but our method is compared with a method with random subgoals. The method with random subgoals is similar to HER.

Hierarchical reinforcement learning(HRL) utilizes subgoals. The option framework is major in the field of HRL. The framework of Sutton *et al.* [16] was able to transfer learned policies in an option. An option consists of an initiation set, an intra-option policy, and a termination function. An option expresses a combination of a subtask and a policy for it. The termination function takes on the role of subgoal because it terminates an option and triggers the switching to another option. Recent methods have found good subgoals for a learner simultaneously with policy learning [17], [18]. The differences with our method are whether the policy is over abstract states or not and whether rewards are generated. The framework intends to recycle a learned policy, but our method focuses on improving learning efficiency.

The RGoal architecture by Ichisugi *et al.* [19] deals with recursive subgoals in HRL. The method solves a MDP

with an augmented state space in which each subgoal has a state space. In multitask settings, the method speeds up learning by sharing subroutines between tasks. The objective is different from ours, and the paper suggests that a non-hierarchical RL method is faster than HRL in single task settings. Murata *et al.* [20] used subgoals for accelerating learning in RL. A reward is generated upon the achievement of each subgoal, and there is a  $Q$ -function for each subgoal. The policy selects an action according to the weighted sum of  $Q$ -functions of subgoals and an original  $Q$ -function. This method is applicable for only table-style value functions. Our method can be applied to function-style value functions.

There are a few methods that can be compared with our method. For fair comparison, the objective of the method and the input, which is a subgoal series should in this case, should be the same as our method.

### C. REWARD SHAPING

Reward shaping has been studied actively. Marom and Rosman [21] proposed the reward shaping based on Bayesian methods called belief reward shaping(BRS). The expected value over the estimated probability distribution of an environmental reward over a state and an action is used as shaping rewards. They proved that the policy learned by Q-learning augmented with BRS is consistent to an optimal policy over original MDP.

The landmark-based reward shaping of Demir *et al.* [22] is the closest to our method. The method shapes only rewards on a landmark using a value function. Their study focused on a POMDP environment, and landmarks automatically become abstract states. We focus on a Markov decision process (MDP) environment, and acquire subgoals from human participants.

Reward shaping in HRL has been studied in [23], [24]. Gao and Toni [23] showed that potential-based reward shaping remains policy invariant to the MAX-Q algorithm. Designing potentials every level is laborious work. We use a single high-level value function as a potential, which reduces the design load. Li *et al.* [24] incorporated an advantage function in high-level state-action space into reward shaping. The reward shaping method in [25] utilized subgoals that are automatically discovered with expert trajectories. The potentials generated every subgoal are different.

Potential-based advice is reward shaping for states and actions [26]. The method shapes the  $Q$ -value function directly for a state and an action, and it makes it easy for a human to give advice to an agent regarding whether an action in an arbitrary state is better or not. Subgoals show what ought to be achieved on the trajectory to a goal. We adopted the shaping of a state value function.

Harutyunyan *et al.* [27] has shown that the  $Q$ -values learned by arbitrary rewards can be used for potential-based advice. The method mainly assumes that a teacher negates the agent's selected action. The method uses failures in trial and errors. In contrast, our method uses successes.

## III. PRELIMINARIES

### A. REINFORCEMENT LEARNING

Reinforcement learning is a framework for acquiring a policy maximizing future rewards through interactions between an agent and an environment. This works under a Markov decision process(MDP). A MDP is represented as a tuple  $(S, A, T, \gamma, R)$ , where  $S$  is a finite set of states,  $A$  is a set of actions of  $k \geq 2$ ,  $T = P(s_{t+1}|s_t, a)$  is a state transition function,  $\gamma \in (0, 1]$  is the discount factor, and  $R$  specifies a reward function. An agent has a policy  $\pi(a|s)$  and a value function  $V(s)$ ,  $Q(s, a)$ . There are two methods for learning an optimal policy, the value-based method [28] and the policy gradient-based method [29]. The value-based method indirectly learns a policy by estimating the optimal value function. The policy gradient-based method directly learns the parameters of a policy.

### B. POTENTIAL BASED REWARD SHAPING

Potential based reward shaping (PBRS) [7], [30] is an effective method for keeping an original optimal policy in an environment with an additional reward function. If the potential-based shaping function  $F$  is formed as

$$F(s_t, s_{t+1}) = \gamma \Phi(s_{t+1}) - \Phi(s_t) \quad (1)$$

it is guaranteed that policies in MDP  $M = (S, A, T, \gamma, R)$  are consistent with those in MDP  $M' = (S, A, T, \gamma, R+F)$ .  $\Phi(s_t)$  is a function with a state  $s_t$  as an argument. Wiewiora showed that a reinforcement learner with initial  $Q$ -values based on the potential function make the same updates throughout learning as a learner receiving potential based shaping reward [30]. We assume two learners  $L$  and  $L'$ .  $L$  initializes the  $Q$ -value to  $Q_0(s, a)$  and updates its policy with an environmental reward  $r$  and the shaping reward  $F$ . The  $Q$ -value of  $L'$  is initialized to  $Q'_0(s, a) = Q_0(s, a) + \Phi(s)$ , and  $L'$  uses only the reward  $r$ . The  $Q$ -values are disassembled into the initial values  $Q_0$  and  $Q'_0$  and the updated values  $\Delta Q(s, a)$  and  $\Delta Q'(s, a)$  through learning, respectively. If  $L$  and  $L'$  learn on the same sequence of experiences,  $\Delta Q(s, a)$  is always equal to  $\Delta Q'(s, a)$ . Moreover, value-based policies of  $L$  and  $L'$  have an identical probability distribution for their next action. Our proposed method uses potential-based shaping rewards to keep an optimal policy in an original MDP.

## IV. SUBGOAL-BASED REWARD SHAPING

We propose subgoal-based reward shaping(SRS), which incorporates human knowledge of subgoals into an algorithm via reward shaping. In [7], it was mentioned that learning sped up with a learned value function used as the potential function. Since a policy learned with potential-based reward shaping is equivalent to that with  $Q$ -value initialization with the potential function [30], using the learned value function for the potential function is equivalent to initializing the value function with the learned value function. It is impossible to prepare the learned values before learning. Therefore, we consider an approximation of the optimal values by using subgoals without a learned value function. To this end,




FIGURE 1. Ordered subgoals.

we define the subgoal first, and we extend the potential-based reward shaping, and then we propose a subgoal-based potential in this section.

**A. SUBGOAL**

We define a *subgoal* as a state  $s$  if  $s$  is a goal in one of the sub-tasks decomposed from a task. In the option framework, the subgoal is the goal of a sub-task, and it is expressed as a termination function [16]. Many studies on the option framework have developed automatic subgoal discovery [17]. We aim to incorporate human subgoal knowledge into the reinforcement learning algorithm with less human effort required. The property of a subgoal might be a part of the optimal trajectories because a human should decompose a task to achieve a goal. We acquire a subgoal series and incorporate subgoals into our method in the experiments. The subgoal series is written formally as  $(SG, <)$ .  $SG$  is a set of subgoals and a sub-set of  $S$ . There are two types of subgoal series, totally ordered and partially ordered as shown as Fig. 1.

With totally ordered subgoals, a subgoal series is deterministically determined at any subgoal. In contrast, partially ordered subgoals have several transitions to the subgoal series from a subgoal. We used only the totally ordered subgoal series in this paper, but both types of ordered subgoals can be used with our proposed reward shaping. Since an agent needs to achieve a subgoal only once, the transition between subgoals is unidirectional.

**B. POTENTIAL BASED REWARD SHAPING WITH STATE HISTORY GUARANTEES THE POLICY INVARIANCE**

We extended potential-based reward shaping to allow the history of the states an agent visited as an input of a potential function. Informally, if the potential function obtains the history of states in the difference equation, the guarantees of policy invariance remain. Formally,

$$F(h_t, h_{t+1}) = \gamma \Phi(h_{t+1}) - \Phi(h_t) \tag{2}$$

where  $h_t$  is the history of the states that an agent visited until state  $s_t$ , expressed as an equation,  $h_t = [s_0, \dots, s_t]$ . To prove the policy invariance, we first define the return in any arbitrary policy  $\pi$  when visiting a state  $s_k$  in a discount manner without shaping. Formally,

$$U_\pi(s_k) = \sum_{t=0}^{\infty} \gamma^t r_{t,\pi} \tag{3}$$

Note that  $t$  is a time step. The Q function can be rewritten as follows:

$$Q_\pi(s_k, a) = \sum_{s_{k+1}} P(s_{k+1}|s_k, a) U_\pi(s_{k+1}) \tag{4}$$

Now consider the same policy but with a reward function modified by adding a potential-based reward function of the form given in Equation (2). The return of the shaped policy  $U_{\pi,\Phi}$  experiencing the same sequence  $s_{k+1}$  is,

$$\begin{aligned} U_{\pi,\Phi}(h_{k+1}) &= \sum_{t=0}^{\infty} \gamma^t (r_{t,\pi} + F(h_{t+k}, h_{t+k+1})) \\ &= \sum_{t=0}^{\infty} \gamma^t (r_{t,\pi} + \gamma \Phi(h_{t+k+1}) - \Phi(h_{t+k})) \\ &= \sum_{t=0}^{\infty} \gamma^t r_{t,\pi} - \Phi(h_k) \\ &= U_\pi(s_k) - \Phi(h_k) \end{aligned} \tag{5}$$

where  $h_k = [s_0, \dots, s_k]$  is written.  $Q_{\pi,\Phi}$  follows from Equation (4) and is expressed as follows.

$$\begin{aligned} Q_{\pi,\Phi}(s_k, a) &= \sum_{s_{k+1}} P(s_{k+1}|s_k, a) U_{\pi,\Phi}(h_{k+1}) \\ &= \sum_{s_{k+1}} P(s_{k+1}|s_k, a) (U_\pi(s_{k+1}) - \Phi(h_k)) \\ &= \sum_{s_{k+1}} P(s_{k+1}|s_k, a) U_\pi(s_{k+1}) \\ &\quad - \sum_{s_{k+1}} P(s_{k+1}|s_k, a) \Phi(h_k) \\ &= Q_\pi(s_k, a) - \Phi(h_k) \end{aligned} \tag{6}$$

Given by Equation (6),  $\Phi$  does not have action  $a$  in parameters. As  $\Phi$  is independent of the action taken, in any given state, the best action remains constant regardless of shaping. Therefore, we can conclude that the guarantee of policy invariance remains.

**C. SUBGOAL-BASED POTENTIALS**

The subgoal-based potential is the main part of our method. We design potentials that utilize subgoals to approximate of the optimal value functions. Generally, the larger the optimal value is, the closer a state is in an environment with only goal reward. We consider subgoal achievement as approaching the goal, and we design a potential that grows every subgoal




FIGURE 2. Potential function for subgoals.

achievement. Formally, we define subgoal-based potentials as

$$\Phi(h) = \eta * c(h) \quad (7)$$

where  $\eta$  is a hyper parameter, and  $c(h)$  is a function for calculating the number of achieved subgoals in a visited state sequence of  $h$  from the last achieved subgoal. An output sequence of the potential function is shown in Fig. 2.

Fig. 2 shows the sequence of an agent achieving three subgoals. When the agent achieves  $sg_1$ , the potential function generates  $\eta * c(h)$ . Upon the achievement of  $sg_2$ , this value becomes larger than  $sg_1$  because  $c(h)$  increases.

#### D. WHOLE ALGORITHM

The pseudocode is shown in Algorithm 1:

---

#### Algorithm 1 Subgoal Potential-Based Reward Shaping

---

**Ensure:**  $\gamma, \eta, SG$

- 1:  $s \leftarrow s_0, h \leftarrow [], h' \leftarrow []$
  - 2: **repeat**
  - 3:   Choose  $a$  according to  $\pi$
  - 4:   Take  $a$  in  $s$ , observe  $s', r$
  - 5:    $h' \leftarrow h$
  - 6:   Append  $s'$  into  $h'$
  - 7:    $\Phi(h') = \eta * c(h')$
  - 8:    $F(h, h') = \gamma \Phi(h') - \Phi(h)$
  - 9:   Update value function and policy with  $F(h, h')$
  - 10:    $s \leftarrow s', h \leftarrow h'$
  - 11: **until** termination.
- 

Though we consider totally ordered subgoals in the experiments, we can easily extend the algorithm with a sequence of partially ordered subgoals. We modify only the function  $c(h)$  for partially ordered subgoals. After an agent takes an action  $a$  and receives a state  $s'$  from an environment, the state  $s'$  is added to the history of states  $h'$ . The calculation of  $\Phi$  involves  $h$  and  $h'$ . The function  $\Phi(h)$  is used for computing  $F(h, h')$ . Even if we only give subgoals without a sequence, the algorithm can work. Since they are partially ordered subgoals that have only a single tier of a sequence of subgoals, the extension mentioned above is available.

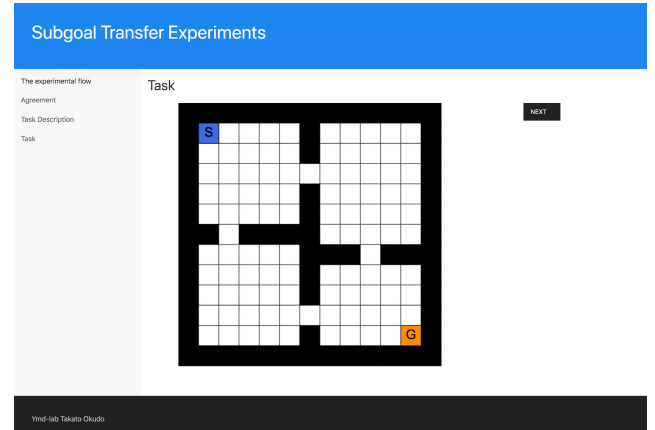


FIGURE 3. UI for acquiring subgoals.

#### V. USER STUDY: HUMAN SUBGOAL ACQUISITION

In this section, we explain a user study done to acquire human knowledge of subgoals. We used navigation tasks in two domains, four-rooms and pinball, that express the state space spatially. Additionally, we conducted a pick and place task, which is a basic robotics task and one for which humans can have difficulty controlling the robot.

##### A. NAVIGATION TASK

An online user study done to acquire human subgoal knowledge using a web-based GUI in Fig. 3 was conducted.

We recruited 10 participants who consisted of half graduate students in the department of computer science and half others (6 males and 4 females, ages 23 to 60, average of 36.4). We confirmed they did not have expertise on subgoals in the two domains. Participants were given the same instructions as follows for the two domains, and they were then asked to designate their two subgoals both for the four-rooms and pinball domains in this fixed order. The number of subgoals was the same as the hallways in the optimal trajectory for the four-rooms domain. The instructions explained to the participants what the subgoals were and how to set them through the GUI shown in Fig. 3. Also, specific explanations of the two task domains were given to the participants. In this experiment, we acquired just two subgoals for learning since they are intuitively considered easy to give on the basis of the structure of the problems. We considered the two subgoals to be totally ordered ones.

##### B. PICK AND PLACE TASK

A user study for the pick and place task was also done online. Since it was difficult to acquire human subgoal knowledge with a GUI, we used a descriptive answer-type form. We assumed that humans use subgoals when they teach behavior in a verbal fashion. They state not how to move but what to achieve in the middle of behavior. The results of this paper minorly support this assumption. We recruited five participants who were amateurs in the field of computer science (3 males and 2 females, ages 23 to 61, average

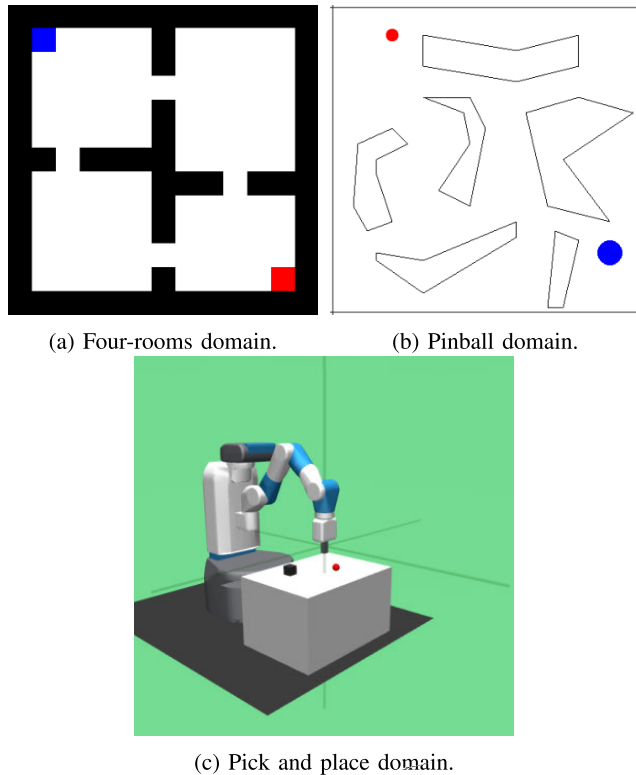


FIGURE 4. Three domains.

of 38.4). The participants read the instructions and then typed the answer in a web form. The instructions consisted of a description of the pick and place task, a movie of manipulator failures, a glossary, and a question on how a human can teach successful behavior. The question included the sentence “Please teach behavior like you would teach your child.” This is because some participants answered that they did not know how to teach the robot in a preliminary experiment. We imposed no limit on the number of subgoals.

## VI. EXPERIMENTS FOR EVALUATION

As mentioned above, we conducted experiments to evaluate our proposed method in three domains: four-rooms, pinball, and pick and place. The four-rooms and pinball domains are navigation tasks in which an agent travels to its goal with the shortest number of steps. The pick and place domain is a robotics task. A robot aims to pick up an object and bring it to a target space. Fig. 4 shows images of the three domains.

We compared human subgoal-based reward shaping (HSRS) with three other methods. They were a baseline RL algorithm, random subgoal potential-based reward shaping (RSRS), and a naive subgoal reward method (NRS). As mentioned in Section II, there are a few methods that can be compared with our method. For fair comparison, the objective of the method needs to be the same as well as the input, that is, a subgoal series. RSRS is a similar approach to the HER algorithm. NRS is a simple method, and many algorithm designers should come up this approach first. In NRS, the function  $\Phi(s)$  in potential-based reward shaping generates a scalar value  $\eta$  just when an agent visits a subgoal

state. The potential function is written formally as follows.

$$\Phi(s) = \begin{cases} \eta & s = sg \\ 0 & s \neq sg \end{cases} \quad (8)$$

The difference from our method is that the positive potential is only for subgoals. A separate baseline RL algorithm was implemented for each domain. This is to show that our method can be adapted to various RL algorithms. The different baseline RL algorithms are used each domain to show the availability for several RL algorithms. The reward shaping method, HSRS, RSRS, and NRS were implemented in the baseline algorithm. HSRS used the participants’ subgoals described in Section V.

Each domain had 20 subgoals the same as the number of participants’ subgoals. For pick and place, we generated two random observations with ranges as subgoals. The subgoals of pinball and pick and place had these ranges because both domains had a continuous state space. We evaluated the learning efficiency with the time to threshold and the asymptotic performance [31] in terms of human subgoal knowledge effects. The time to threshold was the number of episodes to get below a predefined number of threshold steps. The asymptotic performance was the final performance of learning. All the experiments were conducted with a PC [Ryzen 9 5950X, 16 cores(3.4GHz), 128 GB of memory].

### A. NAVIGATION IN FOUR-ROOMS DOMAIN

The four-rooms domain had four rooms, and the rooms were connected by four hallways. The four-rooms domain is a common RL task. In this experiment, learning consisted of 1000 episodes. An episode was a trial run until an agent reached a goal state successfully or when 1000 state actions ended in failure. A state was expressed as a scalar value labeled through all states. An agent could select an action from up, down, left, and right. The agent’s action failed with a probability of 0.33, and another random action was then performed. A reward of +1 was generated when the agent reached a goal state. The start state and goal state were placed at a fixed position. The agent repeated learning 100 times.

#### 1) ALGORITHM SETUP

SARSA is a basic RL algorithm. We used SARSA as a baseline RL algorithm. The value function was in table form. The policy was soft-max. We set the learning rate to 0.01 and the discount rate to 0.99. The function  $c(h)$  counted the number of subgoals in a sequence in order from the front of a state sequence  $h$ . We set  $\eta$  to 0.01 for HSRS, RSRS and NRS after grid search was performed on grids of 0.01, 0.1, 1, 10, and 100.

#### 2) SUBGOALS

Fig. 5 shows the subgoal distribution acquired from the participants and the random subgoal distribution in the four-rooms domain.

The numbers in the cells stand for the frequencies of the subgoals in Fig. 5; higher values are in yellow and lower values are in orange. The number in a cell was the frequency




FIGURE 5. Subgoal distributions of a four-rooms domain.




FIGURE 6. Learning curves compared with three methods.

at which the participants selected that cell as a subgoal. These subgoals are totally ordered subgoals. As shown in Fig. 5(a), participants tended to set the subgoals in the hallways. In Fig. 5(a), the color of the start point, the goal, and subgoals are blue, red, and green, respectively. HSRS and NRS used the subgoals as shown in Fig. 5(a), and RSRS used as shown in Fig. 5(b).

3) RESULTS

We show the results of the learning experiment. Fig. 6 shows the learning curves of our proposed method and the three other methods.

We plotted HSRS with an average totaling 1000 learnings over all participants. RSRS was also averaged by 1000 learnings over all participants. SARSA was averaged by 100 learnings because it does not use subgoals. NRS had almost the same conditions as HSRS. HSRS seemed to have the fewest steps for almost all episodes. The performance of RSRS was close to HSRS but worse than HSRS. NRS had the worst performance. This shows the difficulty of transformation from subgoals into an additional reward function. We also performed an ANOVA for the time to reach the threshold among HSRS, RSRS, SARSA, and NRS. The analysis showed a significant difference at the level of  $p < .01$  among these methods. The Holm-Bonferroni method was used for sub-effect tests. We set the thresholds to 500, 300, 100, and 50 steps. Table 1 shows the mean episodes and the standard deviations of the compared methods for each threshold step.

TABLE 1. Mean and standard deviation:mean(S.D.).

Thres.	HSRS	RSRS	SARSA	NRS
500	2.17(1.54)	3.12(2.43)	3.2(2.61)	6.45(45.1)
300	4.27(3.07)	5.48(3.73)	5.69(3.67)	35.3(156)
100	18.8(9.42)	18.7(10.1)	27.7(11.9)	175(344)
50	35.9(11.7)	37.7(12.4)	53.9(17.8)	376(444)

TABLE 2. Summary of ANOVA and sub-effect tests for time (episode) to threshold steps. \* stands for  $p < .05$ .

Thres.	HSRS<NRS	RSRS<NRS	SARSA<NRS	otherwise
500	*	*	*	n.s.
300	*	*	*	n.s.
100	*	*	*	n.s.
50	*	*	*	n.s.

TABLE 3. Asymptotic performances in four-rooms domain. Ind. is indicator, and S.D. is standard deviation.

Ind.	HSRS	RSRS	SARSA	NRS
Mean	19.1	19.1	24.9	552
S.D.	0.40	0.41	13.5	455




FIGURE 7. Asymptotic performance in four-rooms domain.

Table 2 shows the results of the sub effect tests by the Holm-Bonferroni method for HSRS, RSRS, SARSA and NRS.

As shown in Tables 1 and 2, HSRS, RSRS, and SARSA were significantly lower than NRS but did not have significant differences between the three of them. From Table 1, HSRS shortened the number of learnings by up to 18 episodes from SARSA. Table 6 shows the asymptotic performances of all four methods, and Fig. 7 shows them by box plot.

There were no significant differences among the four methods in terms of asymptotic performance.

B. NAVIGATION IN PINBALL DOMAIN

The navigation task in the pinball domain involves moving a ball to a designated target by giving it velocity. The pinball domain makes it difficult for humans to control the ball because it needs them to control the ball with delicate actions. Delicate control is often necessary in the control domain. It is easier to give a subgoal than trajectories in this domain.

The difference with the four-rooms domain is the continuous state space over the position and velocity of the ball on the x-y plane. An action space has five discrete actions,

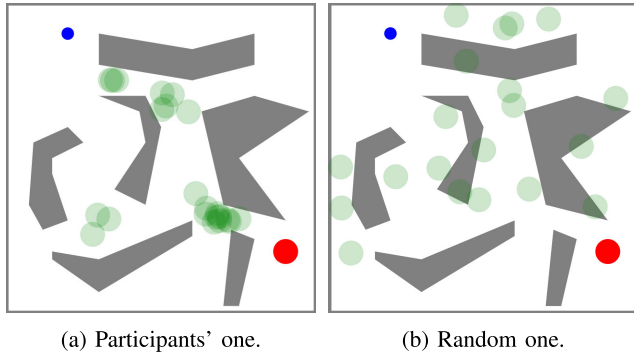


FIGURE 8. Subgoal distributions of a pinball domain.

four types of force and no force. In this domain, a drag coefficient of 0.995 effectively stops ball movements after a finite number of steps when the no-force action is chosen repeatedly; collisions with obstacles are elastic. The four types of force are up, down, right, left on a plane. An episode terminates with a reward of +10,000 when the agent reaches the goal. Interruption of any episode occurs when the episode takes more than 10000 steps. Learnings of the methods are repeated 100 times from scratch, and we used the results to evaluate the learning efficiency.

1) ALGORITHM SETUP

We used the actor-critic algorithm(AC) [32] as the baseline RL algorithm. AC is a basic RL algorithm that directly optimizes a policy with parameters using the value function. A policy with parameters is called the actor part, and a value function is the critic part. Both parts needed to update parameters. We used AC as the critic with linear function approximation over a Fourier basis [33] of order 3. A subgoal had only a center position and a radius in this domain. The radius was the same as that of the target. We assumed that a subgoal was achieved when the ball entered a circle with the center position and radius at any velocity. We used a soft-max policy as the actor. The learning rates were set to 0.01 for both the actor and the critic. The discount factor was set to 0.99. The parameters for HSRS, RSRS, NRS, and  $\eta$  were set to 100 after a grid search on grids of 10, 100, 1,000, and 10,000. The function  $c(h)$  was almost the same as in the four-rooms domain excluding the judgement of whether a state is a subgoal. We compared HSRS with the three other methods in terms of learning efficiency with the time to threshold. We defined this threshold in this domain as the number of episodes required to reach the steps.

2) SUBGOALS

Fig. 8 shows the subgoal distribution acquired from the participants and the random subgoal distribution in the pinball domain.

The participants focused on four regions of branch points to set subgoals in comparison with random subgoals from Fig. 8(a) and Fig. 8(b). HSRS and NRS used the subgoals as shown in Fig. 8(a), and RSRS used those as shown in Fig. 8(b).




FIGURE 9. Learning curve in pinball domain.

TABLE 4. Mean and standard deviation: mean(SD).

Thres.	HSRS	RSRS	AC	NRS
3000	15.0 (32.7)	23.8 (45.7)	51.7 (67.1)	39.8 (50.2)
2000	28.0 (38.8)	33.2 (46.8)	65.4 (67.6)	59.0 (58.6)
1000	63.7 (60.0)	72.4 (72.8)	101 (71.3)	92.0 (69.2)
500	140 (67.0)	141 (14.5)	163 (54.8)	158 (61.0)

TABLE 5. Summary of ANOVA and sub effect tests for time(episode) to threshold steps. Otherwise is abbreviated to o.w., and n.s. means not significant. \* stands for  $p < .05$ .

Thres.	{HSRS, RSRS}<AC	HSRS<NRS	RSRS<NRS	o.w.
3000	*	*	*	n.s.
2000	*	*	*	n.s.
1000	*	*	n.s.	n.s.
500	n.s.	n.s.	n.s.	n.s.

3) RESULTS

Fig. 9 shows the learning curves. It took an average shift of 10 episodes.

HSRS mostly performed the best of the three methods. The performance of RSRS was close to HSRS but worse than HSRS. We evaluated the learning efficiency by using the time to threshold. We used each learning result smoothed by using a simple moving average with the number of time periods being 10 episodes, and we performed a Student’s t-test to confirm the difference among the three methods. Table 4 shows the mean and standard deviations, and Table 5 summarizes the results for the time to threshold, where Thres. is the number of steps. From the result of the ANOVA, there was a significant difference at the level of  $p < .01$  among the four methods for 1000, 2000, and 3000, and there was a difference at  $p < .05$  for 500. The Holm-Bonferroni method was used for the sub-effect tests.


From Table 5, HSRS and RSRS had fewer steps than AC for the thresholds of 3000, 2000, and 1000, which was significant. There was no significant difference between HSRS and RSRS for any of the thresholds. Table 6 shows the asymptotic performances of all four methods, and Fig. 10 shows them by box plot.

There were no significant differences among the four methods in terms of asymptotic performance. Our method with random and human subgoals lead to more efficient learning than the basic RL algorithm in the pinball domain.



**TABLE 6. Asymptotic performances in pinball domain. Ind. is indicator, and S.D. is standard deviation.**

Ind.	HSRS	RSRS	AC	NRS
Mean	3008	3629	3625	3672
S.D.	2678	2925	3014	3042




**FIGURE 10. Asymptotic performance in pinball domain.**

### C. PICK AND PLACE

We used a fetch environment based on the 7-DoF Fetch robotics arm of OpenAI Gym [34]. In pick and place, the robotic arm learns to grasp a box and move it to a target position [35]. We converted the original task into a single-goal reinforcement learning framework because potential-based reward shaping is not developed to deal with the multiple goals [7]. The dimension of observation is larger than the previous navigation tasks, and the action is continuous. An observation is 25-dimensional, and it includes the Cartesian position of the gripper and the object as well as the object’s position relative to the gripper. The reward function generates a reward of  $-1$  every step and a reward of  $0$  when the task is successful. In [35], the task is written about in detail.


#### 1) ALGORITHM SETUP

We compared HSRS with RSRS, NRS, and DDPG [36] in terms of learning efficiency with the time to threshold and asymptotic performance. HSRS, RSRS, and NRS used DDPG as a base. We defined this threshold in the task as the number of epochs required to reach the designated success rate. The asymptotic performance was the average success rate between 190 and 200 epochs. Ten workers stored episodes and calculated the gradient simultaneously in an epoch. HSRS and NRS used ordered subgoals provided by five participants. RSRS used subgoals randomly generated. The learning in 200 epochs took several hours. We used an OpenAI Baselines [37] implementation for DDPG with default hyper-parameter settings. We built the hidden and output layers of the value network over abstract states with the same structure as the Q-value network. We excluded the action from the input layer. The input of the network was only the observation of subgoal achievement, and the network learned from the discount accumulated reward until subgoal achievement. A subgoal was defined from the information in the observation. We set the margin to  $\pm 0.01$  to loosen the



(a) First subgoal. (b) Second subgoal.

**FIGURE 11. Example of subgoals from participant in a pick and place domain.**



**FIGURE 12. Learning curves in pick and place task.**

severe condition to achieve subgoals.  $\eta$  was set as 1 for HSRS, RSRS, and NRS. The learning repeated 10 times, and the results were averaged over 10 learnings.

#### 2) SUBGOALS

All five participants determined the subgoal series, the first subgoal was reaching the location available to grasp the object, and the second subgoal was grasping the object. Fig. 11 shows the example of the subgoal series in pick and place domain.

#### 3) RESULTS

We used the subgoal series for the input of our method. Fig. 12 shows the learning curves of HSRS, RSRS, NRS, and DDPG.

As shown in Fig. 12, the results were averaged across five learnings, and the shaded areas represent one standard error. The random seeds and the locations of the goal and object were varied every learning. HSRS worked more effectively than DDPG, especially after about the 25th epoch. The performances of RSRS and NRS were similar. NRS had the worst performance through almost all epochs.

Table 7 shows the mean episodes and the standard deviations of the compared methods for each threshold steps.

Table 8 shows the asymptotic performances of all four methods, and Fig. 13 shows them by box plot.


HSRS was significantly different from RSRS and NRS at 0.6 in terms of the time to threshold. For the rest, there were no differences in both the time to threshold and the asymptotic

**TABLE 7. Mean and standard deviation in pick and place:mean(S.D.).**

Thres.	HSRS	RSRS	DDPG	NRS
0.2	26.3 (22.3)	27.1 (6.77)	28.2 (26.4)	34.3 (18.7)
0.4	42.7 (27.2)	87.0 (49.6)	56.6 (37.5)	75.9 (43.5)
0.6	75.1 (56.7)	175 (46.4)	110 (63.7)	189 (26.2)
0.8	156 (57.6)	185 (35.7)	187 (22.9)	200 (0)

**TABLE 8. Asymptotic performances in pick and place. Ind. is indicator, and S.D. is standard deviation.**

	HSRS	RSRS	DDPG	NRS
Mean	0.73	0.53	0.63	0.51
S.D.	0.16	0.17	0.13	0.06



**FIGURE 13. Asymptotic performance in pick and place task.**


performance, but HSRS was the best compared with the other three methods.

In this section, we evaluated our method for four-rooms, pinball, and pick and place. In all domains, HSRS performed better than the other three methods in terms of the time to threshold and the asymptotic performance. There was, in particular, a significant difference between HSRS and AC as the baseline RL algorithm for pinball. The performance of RSRS was close to HSRS in the four-rooms and pinball domains. In contrast, RSRS performed the second worst of all for pick and place. The results showed the subgoal-based reward shaping with human knowledge of subgoal improved the learning efficiency from the baseline RL algorithm, and the human knowledge was more useful for SRS than random. The naive reward shaping made the learning efficiency worse than the baseline RL algorithm, so it is not easy for designers to incorporate the knowledge of subgoal over PBRs. This is why we proposed SRS, and SRS could improve the learning efficiency just by providing the subgoal series.


**VII. DISCUSSION**

**A. ANALYZING PERFORMANCES**

There was a small difference between HSRS and RSRS in the four-rooms domain from Fig. 6 and the pinball domain from Fig. 9. The difference between HSRS and RSRS was large in the pick and place domain. Approximately 65% of states generated randomly were in an optimal trajectory for the four-rooms domain. The pinball task seemed to have approximately 20% of states generated randomly in the optimal trajectory. There was no state generated randomly in an



(a) Four-rooms domain.



(b) Pinball domain.

**FIGURE 14. Results of grid searches.**

optimal trajectory in the pick and place task. The random subgoals were better in the four-rooms and pinball domains than in the pick and place task. This is because these two domains might have had more states in the optimal trajectory than the pick and place task. We think that the small difference between HSRS and RSRS was caused by the domains having states in the optimal trajectory. The pick and place task has a large state space, so the random subgoals are not likely to be placed on the optimal trajectories. The larger the state space becomes, the larger the difference between HSRS and RSRS may become.

Potential-based reward shaping keeps a policy invariant from the transformation of a reward function. From the experimental results in the pinball domain, the asymptotic performances of HSRS were statistically significantly lower than RSRS. There was no significant difference in the four-rooms domain. As shown in Fig. 6, the performance was clearly asymptotic at the 121th of 1000 episodes. The learning was clearly asymptotic as the 200th episode for the pinball domain in Fig. 9. Since our proposed method is based on PBRs, RSRS converges to the same performance as HSRS if learning continues.

**B. TUNING HYPER PARAMETERS**

Fig. 14 shows the results of grid searches. In the searches, we used only a single participant’s subgoals which were

randomly selected, and the number of learnings was a half that of the experimental results.

As seen in this figure, our method was significantly sensitive to the hyper parameter  $\eta$ . A wrong  $\eta$  such as  $\eta = 1$  in the four-rooms domain could deteriorate the performance of the method. This may incur a cost for tuning hyper parameters. To solve this problem, we must fit  $\eta$  to the optimal value. The environmental rewards of goals were 1 and 10,000 for four-rooms and pinball, respectively. The best  $\eta$  was 0.01 for four-rooms and 100 for pinball. Both  $\eta$  were one-hundredth of the environmental rewards.

### C. LIMITATIONS

#### 1) DIFFICULTIES OF HYPER PARAMETERS

Since we didn't know the preferable hyper parameter  $\eta$  explicitly, we tuned hyper parameters by grid searches. The traditional reward shaping method has tackled the same problem. The SARSA-RS method uses a value function over abstract states as a potential function [38]. The method doesn't require hyper parameters of the potential function. If we can make SARSA-RS incorporate a series of subgoals, we can ignore the hyper parameter  $\eta$  because  $\eta$  can be obtained through learning. To do this, we proposed dynamic trajectory aggregation in another paper [39]. The method uses a sequence of subgoals to build abstract states. The part of a trajectory until subgoal achievement becomes an abstract state, and the part until the next subgoal achievement becomes another abstract state. The value of the abstract states is used for shaping. With the method, the designer provides only a sequence of subgoals.

#### 2) TEACHING KNOWLEDGE OF SUBGOALS

Our method depends on whether the teacher can provide the sequence of subgoals in a task, and can be used to the task which the teacher can provide it. The limitation of subgoals is that no established method to provide them, and it may depend on individual intuitive consideration. For the four-rooms domain, as shown in Fig. 5, almost all the subgoals were located within the right-top and right-bottom rooms. From this, we think that many participants tended to consider the right-down path as the shortest one. This may mean humans abstractly have a common methodology and preference for giving subgoals. Additionally, there was the interesting observation that half of the participants set a subgoal in the hallways.

For the pinball domain, as shown in Fig. 8, the subgoals were scattered at the gap between objects and at the branch points. They seem to be at the change points of actions on an optimal trajectory. We can consider that the repeated actions are composed in a macro-action. This is a similar approach to [40].

For the pick and place task, five participants provided the same subgoals. The way to provide subgoals was different from the other tasks. We used the descriptive answer-type form. These results almost show the tendency for people to provide similar subgoals exists.

We consider that people change a strategy for providing subgoals in response to the task properties, especially environment structure. The test for this hypothesis is future work.

#### 3) DIFFICULTIES OF PROVIDING SUBGOALS

There are other limitations and future work regarding this work. We implicitly assumed that a human can set subgoals more easily than providing the optimal trajectories from a start state to a goal state, which is conventional human knowledge that improves efficiency of reinforcement learning with IRL. We need to verify this assumption by conducting experiments with participants to measure the cognitive load in setting various subgoals and trajectories.

The ease of providing subgoals may depend on the kind of task. We think that there are two steps that a human can take to set subgoals in a task. In the first, the human comes up with an optimal trajectory for a task. The optimal trajectory includes subgoals. If the human cannot come up with this trajectory, he or she cannot know the subgoals. In the second, the human selects states as subgoals in a trajectory. The trajectory includes states, and the human must prioritize them because the high-prioritized states become subgoals. A task in which a human can process these two steps easily would be an appropriate one for subgoals.

We plan to conduct a large user study to make clear what the best task for providing subgoals is. The viewpoint of the human teacher may be a factor. For this, we can consider a bird's-eye view and a first-person view as levels. The navigation tasks in this paper involved a bird's-eye view. An example of a task with a first-person view is a first-person shooter game such as ViZDoom [41]. Our hypothesis is that a task with a bird's-eye view is more suitable for providing subgoals than one with a first-person view. Another factor may be the explainability of the optimal trajectories. The trajectory of the four-rooms domain is easy for a human teacher to explain, such as explaining that the agent should first go the top-right hallway and then the bottom-right hallway. The teacher can select subgoals easily in this domain like the top-right and bottom-right hallways. In the Hand Manipulation task [35], it is hard for the teacher to explain how the robotic hand of the agent should move. It may be also difficult for the teacher to select subgoals. These are related to explicit and implicit knowledge, respectively. Our hypothesis is that a task including an optimal trajectory that the teacher can explain is more suitable than one in which the teacher cannot explain the trajectory under subgoal settings. Future work is also to test these hypotheses in a large user study.

### VIII. CONCLUSION

In reinforcement learning, learning a policy is time-consuming. We aim for accelerating learning with reward transformation based on human subgoal knowledge. The difficulty of integrating subgoal knowledge into the potentials was a problem in improving learning efficiency. We proposed a method by which a human deals with several characteristic states as a subgoal. We defined a subgoal as the goal state

in one of the sub-tasks into which a human decomposes a task. The main part of our method is an approximation of the optimal value function with subgoals. We collected ordered subgoals from participants and used them for evaluation. We evaluated navigation for four-rooms, pinball, and a pick and place task. The experimental results revealed that our method with human subgoals enabled faster learning compared with the baseline method, and human subgoal series were more helpful than random ones. The availability of our method depends on the ease of getting subgoals. The limitations of subgoals are that there is no established methodology for giving them and what kinds of tasks are suitable for them is unknown. Elucidating this is our future work.

## REFERENCES

- [1] M. E. Taylor, "Improving reinforcement learning with human input," in *Proc. 27th Int. Joint Conf. Artif. Intell.*, Jul. 2018, pp. 5724–5728.
- [2] Z. Wang and M. E. Taylor, "Effective transfer via demonstrations in reinforcement learning: A preliminary study," in *Proc. AAAI Spring Symposia*, Stanford University, Palo Alto, CA, USA, Mar. 2016, pp. 77–83.
- [3] S. Griffith, K. Subramanian, J. Scholz, C. L. Isbell, and A. L. Thomaz, "Policy shaping: Integrating human feedback with reinforcement learning," in *Proc. Adv. Neural Inf. Process. Syst.*, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, Eds. Red Hook, NY, USA: Curran Associates, 2013, pp. 2625–2633.
- [4] Y. Aytar, T. Pfaff, D. Budden, T. L. Paine, Z. Wang, and N. D. Freitas, "Playing hard exploration games by watching youtube," in *Proc. Adv. Neural Inf. Process. Syst.* Red Hook, NY, USA: Curran Associates, 2018, pp. 2935–2945.
- [5] A. Y. Ng and S. J. Russell, "Algorithms for inverse reinforcement learning," in *Proc. 17th Int. Conf. Mach. Learn.*, 2000, pp. 663–670.
- [6] J. Ho and S. Ermon, "Generative adversarial imitation learning," in *Proc. Adv. Neural Inf. Process. Syst.* Red Hook, NY, USA: Curran Associates, 2016, pp. 4565–4573.
- [7] A. Y. Ng, D. Harada, and S. J. Russell, "Policy invariance under reward transformations: Theory and application to reward shaping," in *Proc. 16th Int. Conf. Mach. Learn.*, 1999, pp. 278–287.
- [8] S. Arora and P. Doshi, "A survey of inverse reinforcement learning: Challenges, methods and progress," *Artif. Intell.*, vol. 297, Aug. 2021, Art. no. 103500. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0004370221000515>
- [9] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mané, "Concrete problems in AI safety," 2016, *arXiv:1606.06565*. [Online]. Available: <http://arxiv.org/abs/1606.06565>
- [10] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 3rd ed. Upper Saddle River, NJ, USA: Prentice-Hall, 2009.
- [11] P. F. Christiano, J. Leike, T. Brown, M. Martic, S. Legg, and D. Amodei, "Deep reinforcement learning from human preferences," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 4299–4307.
- [12] S. Ross, G. Gordon, and D. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," in *Proc. 14th Int. Conf. Artif. Intell. Statist.* in Proceedings of Machine Learning Research, vol. 15, G. Gordon, D. Dunson, and M. Dudík, Eds. Fort Lauderdale, FL, USA: PMLR, Apr. 2011, pp. 627–635. [Online]. Available: <http://proceedings.mlr.press/v15/ross11a.html>
- [13] F. Cruz, S. Magg, Y. Nagai, and S. Wermter, "Improving interactive reinforcement learning: What makes a good teacher?" *Connection Sci.*, vol. 30, no. 3, pp. 306–325, Jul. 2018.
- [14] W. B. Knox and P. Stone, "Interactively shaping agents via human reinforcement: The TAMER framework," in *Proc. 5th Int. Conf. Knowl. Capture (K-CAP)*, 2009, pp. 9–16.
- [15] M. Andrychowicz, F. Wolski, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, O. P. Abbeel, and W. Zaremba, "Hindsight experience replay," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 5048–5058.
- [16] R. S. Sutton, D. Precup, and S. Singh, "Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning," *Artif. Intell.*, vol. 112, nos. 1–2, pp. 181–211, Aug. 1999.
- [17] P.-L. Bacon, J. Harb, and D. Precup, "The option-critic architecture," in *Proc. 31st AAAI Conf. Artif. Intell.*, 2017, pp. 1726–1734.
- [18] A. S. Vezhnevets, S. Osindero, T. Schaul, N. Heess, M. Jaderberg, D. Silver, and K. Kavukcuoglu, "Feudal networks for hierarchical reinforcement learning," in *Proc. 34th Int. Conf. Mach. Learn. (ICML)*, vol. 70, Aug. 2017, pp. 3540–3549.
- [19] Y. Ichisugi, N. Takahashi, H. Nakada, and T. Sano, "Hierarchical reinforcement learning with unlimited recursive subroutine calls," in *Artificial Neural Networks and Machine Learning*. Cham, Switzerland: Springer, 2019, pp. 103–114.
- [20] J. Murata, Y. Abe, and K. Ota, "Introduction and control of subgoals in reinforcement learning," in *Proc. 25th Conf. Int. Multi-Conf., Artif. Intell. Appl.* Innsbruck, Austria: ACTA Press, 2007, pp. 329–334.
- [21] O. Marom and B. Rosman, "Belief reward shaping in reinforcement learning," in *Proc. 32nd AAAI Conf. Artif. Intell. (AAAI)*. Menlo Park, CA, USA: AAAI Press, 2018, pp. 3762–3769.
- [22] A. Demir, E. Çilden, and F. Polat, "Landmark based reward shaping in reinforcement learning with hidden states," in *Proc. 18th Int. Conf. Auto. Agents MultiAgent Syst.*, 2019, pp. 1922–1924.
- [23] Y. Gao and F. Toni, "Potential based reward shaping for hierarchical reinforcement learning," in *Proc. 24th Int. Conf. Artif. Intell.* Menlo Park, CA, USA: AAAI Press, 2015, pp. 3504–3510.
- [24] S. Li, R. Wang, M. Tang, and C. Zhang, "Hierarchical reinforcement learning with advantage-based auxiliary rewards," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 32, 2019, pp. 1409–1419.
- [25] S. Paul, J. Vanbaar, and A. Roy-Chowdhury, "Learning from trajectories via subgoal discovery," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 8411–8421.
- [26] E. Wiewiora, G. Cottrell, and C. Elkan, "Principled methods for advising reinforcement learning agents," in *Proc. 20th Int. Conf. Mach. Learn.*, 2003, pp. 792–799.
- [27] A. Harutyunyan, S. Devlin, P. Vrancic, and A. Nowe, "Expressing arbitrary reward functions as potential-based advice," in *Proc. 29th AAAI Conf. Artif. Intell.* Menlo Park, CA, USA: AAAI Press, 2015, pp. 2652–2658.
- [28] C. Szepesvari, *Algorithms for Reinforcement Learning*. San Rafael, CA, USA: Morgan & Claypool, 2010.
- [29] R. S. Sutton, D. A. McAllester, S. P. Singh, and Y. Mansour, "Policy gradient methods for reinforcement learning with function approximation," in *Proc. Adv. Neural Inf. Process. Syst.*, 2000, pp. 1057–1063.
- [30] E. Wiewiora, "Potential-based shaping and Q-value initialization are equivalent," *J. Artif. Intell. Res.*, vol. 19, pp. 205–208, Sep. 2003.
- [31] M. E. Taylor and P. Stone, "Transfer learning for reinforcement learning domains: A survey," *J. Mach. Learn. Res.*, vol. 10, pp. 1633–1685, Jul. 2009.
- [32] V. R. Konda and J. N. Tsitsiklis, "Actor-critic algorithms," in *Proc. Adv. Neural Inf. Process. Syst.*, 2000, pp. 1008–1014.
- [33] G. Konidaris, S. Osentoski, and P. Thomas, "Value function approximation in reinforcement learning using the Fourier basis," in *Proc. 25nd AAAI Conf. Artif. Intell.*, 2011, pp. 380–385.
- [34] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "Openai gym," 2016, *arXiv:1606.01540*. [Online]. Available: <https://arxiv.org/abs/1606.01540>
- [35] M. Plappert, M. Andrychowicz, A. Ray, B. McGrew, B. Baker, G. Powell, J. Schneider, J. Tobin, M. Chociej, P. Welinder, V. Kumar, and W. Zaremba, "Multi-goal reinforcement learning: Challenging robotics environments and request for research," 2018, *arXiv:1802.09464*. [Online]. Available: <https://arxiv.org/abs/1802.09464>
- [36] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," in *Proc. ICLR*, 2016.
- [37] P. Dhariwal, C. Hesse, O. Klimov, A. Nichol, M. Plappert, A. Radford, J. Schulman, S. Sidor, Y. Wu, and P. Zhokhov. (2017). *Openai Baselines*. [Online]. Available: <https://github.com/openai/baselines>
- [38] M. Grzes and D. Kudenko, "Online learning of shaping rewards in reinforcement learning," *Neural Netw.*, vol. 23, no. 4, pp. 541–550, May 2010.
- [39] T. Okudo and S. Yamada, "Reward shaping with dynamic trajectory aggregation," in *Proc. Int. Joint Conf. Neural Netw. (IJCNN)*, Jul. 2021.
- [40] A. S. Lakshminarayanan, S. Sharma, and B. Ravindran, "Dynamic action repetition for deep reinforcement learning," in *Proc. 31st AAAI Conf. Artif. Intell.* Menlo Park, CA, USA: AAAI Press, 2017, pp. 2133–2139.
- [41] M. Kempka, M. Wydmuch, G. Runc, J. Toczek, and W. Jaskowski, "ViZDoom: A doom-based AI research platform for visual reinforcement learning," 2016, *arXiv:1605.02097*. [Online]. Available: <https://arxiv.org/abs/1605.02097>

• • •