

Efficient Exploration in Resource-Restricted Reinforcement Learning

Zhihai Wang^{1*}, Taoxing Pan^{1*}, Qi Zhou¹, Jie Wang^{1,2†}

¹CAS Key Laboratory of Technology in GIPAS, University of Science and Technology of China

²Institute of Artificial Intelligence, Hefei Comprehensive National Science Center

{zhwangx, tx1997,zhouqida}@mail.ustc.edu.cn

jiewangx@ustc.edu.cn

Abstract

In many real-world applications of reinforcement learning (RL), performing actions requires consuming certain types of resources that are non-replenishable in each episode. Typical applications include robotic control with limited energy and video games with consumable items. In tasks with non-replenishable resources, we observe that popular RL methods such as soft actor critic suffer from poor sample efficiency. The major reason is that, they tend to exhaust resources fast and thus the subsequent exploration is severely restricted due to the absence of resources. To address this challenge, we first formalize the aforementioned problem as a resource-restricted reinforcement learning, and then propose a novel **resource-aware exploration bonus** (RAEB) to make reasonable usage of resources. An appealing feature of RAEB is that, it can significantly reduce unnecessary resource-consuming trials while effectively encouraging the agent to explore unvisited states. Experiments demonstrate that the proposed RAEB significantly outperforms state-of-the-art exploration strategies in resource-restricted reinforcement learning environments, improving the sample efficiency by up to an order of magnitude.

1 Introduction

Performing actions requires consuming certain types of resources in many real-world decision-making tasks (Kormushev et al. 2011; Kempka et al. 2016; Bhatia, Varakantham, and Kumar 2019; Cui et al. 2020). For example, the availability of the action “jump” depends on the remaining energy in robotic control (Grizzle et al. 2014). Another example is that the availability of a specific skill depends on a certain type of in-game items in video games (Vinyals et al. 2019). Therefore, making reasonable usage of resources is one of the keys to success in decision-making tasks with limited resources. (Bhatia, Varakantham, and Kumar 2019). In recent years, reinforcement learning (RL) has achieved success in decision-making tasks from games to robotic control in simulators (Schulman et al. 2015; Silver et al. 2017). However, RL with limited resources has not been well studied, and RL methods can struggle to

make reasonable usage of resources. In this paper, we take the first step towards studying RL with resources that are non-replenishable in each episode, which can be extremely challenging and is very common in the real world. Typical applications include robotic control with non-replenishable energy (Kormushev et al. 2011; Kormushev, Calinon, and Caldwell 2013) and video games with non-replenishable in-game items (Kempka et al. 2016; Resnick et al. 2018).

This paper begins by evaluating several popular RL algorithms, including proximal policy optimization (PPO) (Schulman et al. 2017) and soft actor critic (SAC) (Haarnoja et al. 2018), in various tasks with non-replenishable resources. We find these algorithms struggle to explore the environments efficiently and suffer from poor sample efficiency. Moreover, we empirically show that the surprise-based exploration method (Achiam and Sastry 2017), one of the state-of-the-art advanced exploration strategies, still suffers from poor sample efficiency. Even worse, we observe that some of these algorithms struggle to perform better than random agents in these tasks. We further perform an in-depth analysis of this challenge, and find that the exploration is severely restricted by the resources. As the available actions depend on the remaining resources and these algorithms tend to exhaust resources rapidly, the subsequent resource-consuming exploration is severely restricted. However, resource-consuming actions are usually significant for achieving high rewards, such as consuming consumable items in video games. Therefore, these algorithms suffer from inefficient exploration. (See Section 5)

To address this challenge, we first formalize the aforementioned problems as a **resource-restricted reinforcement learning** (R3L), where the available actions largely depend on the remaining resources. Specifically, we augment the Markov Decision Process (Sutton and Barto 2018) with resource-related information that is easily accessible in many real-world tasks, including a key map from a state to its remaining resources. We then propose a novel **resource-aware exploration bonus** (RAEB) that significantly improves the exploration efficiency by making reasonable usage of resources. Based on the observation that the accessible state set of a given state—which the agent can possibly reach from the given state—largely depends on the remaining resources of the given state and large accessible state sets are essential for efficient exploration in these

*Equal contribution.

†Corresponding author.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

tasks, RAEB encourages the agent to explore unvisited states that have large accessible state sets. Specifically, we quantify the RAEB of a given state by its novelty and remaining resources, which simultaneously promotes novelty-seeking and resource-saving exploration.

To compare RAEB with the baselines, we design a range of robotic delivery and autonomous electric robot tasks based on Gym (Brockman et al. 2016) and Mujoco (Todorov, Erez, and Tassa 2012). In these tasks, we regard goods and/or electricity as resources (see Section 5). Experiments demonstrate that the proposed RAEB significantly outperforms state-of-the-art exploration strategies in several challenging R3L environments, improving the sample efficiency by up to an order of magnitude. Moreover, we empirically show that our proposed approach significantly reduces unnecessary resource-consuming trials while effectively encouraging the agent to explore unvisited states.

2 Related Work

Decision-making tasks with limited resources Some work has studied the applications of reinforcement learning in specific resource-related tasks, such as job scheduling (Zhang and Dietterich 1995; Cui et al. 2020), resource allocation (Tesauro et al. 2005), and unpowered glider soaring (Chung, Lawrance, and Sukkarieh 2014). However, they are application-customized, and thus general problems of RL with limited resources have not been well studied. In this paper, we take the first step towards studying RL with resources that are non-replenishable in each episode.

Intrinsic reward-based exploration Efficient exploration remains a major challenge in reinforcement learning. It is common to generate intrinsic rewards to guide efficient exploration. Existing intrinsic reward-based exploration methods include count-based (Strehl and Littman 2004; Bellemare et al. 2016), prediction-error-based (Pathak et al. 2017; Achiam and Sastry 2017), information-gain-based (Houthoof et al. 2016; Shyam, Jaskowski, and Gomez 2019), and empowerment-based (Mohamed and Jimenez Rezende 2015) methods. Our proposed RAEB is also an intrinsic reward. However, RAEB promotes efficient exploration in R3L tasks while previous exploration methods struggle to explore R3L environments efficiently.

Constrained reinforcement learning Constrained reinforcement learning is an active topic in RL research. Constrained reinforcement learning methods (Achiam et al. 2017; Ding et al. 2017; Chen et al. 2020) address the challenge of learning policies that maximize the expected return while satisfying the constraints. Although there exist resource constraints that the quantity of resources is limited in R3L problems, the constraints can be easily introduced to the environment or the agent. For example, we can set the resource-consuming actions unavailable if the resources are exhausted in R3L environments. Thus, any policy can satisfy the resource constraints. In contrast, to find optimal policies in the R3L setting, one of the major challenges lies in efficient exploration.

3 Preliminaries

We introduce the notation we will use throughout the paper. We consider an infinite horizon Markov Decision Process (MDP) denoted by a tuple $(\mathcal{S}, \mathcal{A}, p, r, \gamma)$, where the state space $\mathcal{S} \subset \mathbb{R}^m$ and the action space $\mathcal{A} \subset \mathbb{R}^n$ are continuous, $p : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, \infty)$ is the transition probability distribution, $r : \mathcal{S} \times \mathcal{A} \rightarrow [r_{\min}, r_{\max}]$ is the reward function, $\gamma \in (0, 1)$ is a discount factor. Let the policy π maps each state $s \in \mathcal{S}$ to a probability distribution over the action space \mathcal{A} . That is, $\pi(\cdot|s)$ is the probability density function (PDF) over the action space \mathcal{A} . Define the set of feasible policies as Π . In reinforcement learning, we aim to find the policy π that maximizes the cumulative rewards

$$\pi^* = \arg \max_{\pi \in \Pi} \mathbb{E}_{\pi} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \right], \quad (1)$$

where $s_0 \sim \mu$, $a_t \sim \pi(\cdot|s_t)$, $s_{t+1} \sim p(\cdot|s_t, a_t)$, and μ is the initial distribution of state.

We define the probability density of the next state s' after one step transition following the policy π as

$$p_1^{\pi}(s'|s) = \int_{a \in \mathcal{A}} p(s'|s, a) \pi(a|s) da.$$

Then, we recursively calculate the probability density of the state s' after $t + 1$ steps transition following the policy π by

$$p_{t+1}^{\pi}(s'|s) = \iint_{s_t, a_t} p(s'|s_t, a_t) \pi(a_t|s_t) p_t^{\pi}(s_t|s) da_t ds_t.$$

4 Resource-Restricted Reinforcement Learning (R3L)

We present a detailed formulation of RL with limited resources in this section. We empirically show that the resource-related information is critical for efficient exploration, as the available actions largely depend on the remaining resources (see Section 5). However, the general reinforcement learning formulation assumes the environment is totally unknown and neglects the resource-related information. To tackle this problem, we formalize the tasks with resources as a resource-restricted reinforcement learning, which introduces accessible resource-related information.

4.1 A Detailed Formulation of R3L

We define certain objects that performing actions requires as resources in reinforcement learning settings, such as the energy in robotic control and consumable items in video games. Suppose that there are d types of resources. We use the notation $\xi \in \mathbb{R}_+^d$ to denote the resource vector and $\mathcal{S}_r \subset \mathbb{R}_+^d$ to denote the set of all possible resource vectors.

We augment the state space \mathcal{S} with \mathcal{S}_r . That is, the state space in R3L is $\mathcal{S}_{r3l} = \{s | s = [s_o, s_r], \forall s_o \in \mathcal{S}, \forall s_r \in \mathcal{S}_r\}$, where $[s_o, s_r]$ denotes the concatenation of s_o and s_r . With this augmented state space, algorithms can learn resource-related information from data.




Figure 1: (a) The results show that the average return of Surprise on Electric Mountain Car is much lower than that on Continuous Mountain Car. (b) The results show that the average score of Surprise on Delivery Mountain Car increases with the initial number of goods. (c) We evaluate the learned policies every 10000 training steps on Delivery Mountain Car. The dashed curves correspond to the max height that the car can reach before exhausting the resources, while the solid curves correspond to the max height that the car can reach regardless of whether the car exhausts the resources or not. The results demonstrate that PPO, SAC, and Surprise exhaust the resources when the car only reaches a low height, and then the subsequent exploration is severely restricted due to the absence of resources.

To better exploit the resource-related information, we define a resource-aware function $I : \mathcal{S} \rightarrow \mathbb{R}_+^d$ from the state to the quantity of its available resources,

$$I(s) = (I_1(s), \dots, I_d(s)), I_i(s) \geq 0, \forall i \in [d],$$

where $[d] = \{1, 2, \dots, d\}$. We call the i -th type of resources non-replenishable if the quantity is monotonically non-increasing in an episode, i.e., $I_i(s_t) \geq I_i(s_{t+1})$. We assume I is known a priori because the information about the current resources is easily accessible in many real-world tasks.

For completeness, we further define a deterministic resource transition function $f : \mathcal{S}_{r3l} \times \mathcal{A} \rightarrow \mathcal{S}_r$ which represents the transition function of the resources. We assume f is unknown as the dynamic of the resources is often inaccessible in real-world tasks. Thus, the transition probability distribution is defined by $p_{r3l}(s'|s, a) = p(s'_o | s_o, a) \mathbb{I}(s'_r = f(s, a))$, where $s = [s_o, s_r]$, $s' = [s'_o, s'_r]$, and $\mathbb{I}(\cdot)$ denotes the indicator function. For simplicity, we also denote by p the transition probability distribution in R3L problems.

We call the RL problems with limited resources as R3L problems, whose formulation is denoted by a tuple $(\mathcal{S}_{r3l}, \mathcal{A}, p, r, \gamma, I, f)$. Let $\mathcal{A}(s)$ denote the available action set of a given state s , and we have $\cup_{s \in \mathcal{S}_{r3l}} \mathcal{A}(s) = \mathcal{A}$. Note that the $\mathcal{A}(s)$ depends on the remaining resources at state s . The feasible policy $\pi(\cdot|s)$ is a PDF over the action space $\mathcal{A}(s)$. We also denote by Π the feasible policy set in R3L problems. The same as conventional RL, R3L also aims to solve the problem (1) to find the optimal feasible policy.

Moreover, we call the state s_2 is accessible from the state s_1 , if there is a stationary policy $\pi \in \Pi$ and a time step $t \in \mathbb{N}_+$, such that the probability density $p_t^\pi(s_2|s_1) > 0$. Suppose the resources are non-replenishable and $\exists i \in [d], I_i(s_1) < I_i(s_2)$, then s_2 is inaccessible from s_1 . Besides, we define the set of accessible states of a state s as

$$A_c(s) \doteq \{s' : s' \text{ is accessible from } s\}.$$

5 Challenge in R3L Tasks

In general, we divide real-world decision-making tasks with non-replenishable resources into two categories. In the first category of tasks, all actions consume resources and different actions consume different quantities of resources, such as robotic control with limited energy. In these tasks, the agent needs to seek actions that achieve high rewards while consuming small quantities of resources. In the second category, only specific actions consume resources, such as video games with consumable items. In these tasks, the agent needs to seek proper states to consume the resources.

To evaluate popular RL methods in both kinds of R3L tasks, we design three series of environments with limited resources based on Gym (Brockman et al. 2016) and Mujoco (Todorov, Erez, and Tassa 2012). The first is the autonomous electric robot task. In this task, the resource is electricity and all actions consume electricity. The quantity of consumed electricity depends on the amplitude of the action a , defined by $0.1 * \|a\|_2^2$. The second is the robotic delivery task. In this task, the resource is goods and only the “unload” action consumes the goods. The agent needs to “unload” the goods at an unknown destination. The third is a task that combines the first and the second. Please refer to Appendix B.1 for more details about these environments. Specifically, in this part, we use two designed R3L tasks, namely Electric Mountain Car and Delivery Mountain Car. The two tasks are based on the classic control environment—Continuous Mountain Car. On Electric Mountain Car, the agent aims to reach the mountain top with limited electricity. On Delivery Mountain Car, the agent aims to deliver the goods to the mountain top.

5.1 Inefficient Exploration in R3L Tasks

Based on the designed R3L tasks, we first show that a state-of-the-art exploration method, i.e., the surprise-based exploration method (Surprise) (Achiam and Sastry 2017), suffers from poor sample efficiency in both kinds of R3L tasks, especially those with scarce resources. We compare the performance of Surprise on Electric Mountain Car and Continuous

Mountain Car (without electricity restriction) as shown in Figure 1a. The results show that the performance of Surprise on Electric Mountain Car is much poorer than that on Continuous Mountain Car. We further compare the performance of Surprise on Delivery Mountain Car with different initial quantities of goods. Figure 1b shows that the performance of Surprise improves with the initial quantity of goods. In particular, Surprise struggles to learn a policy better than random agents when resources are scarce.

To provide further insight into the challenge in R3L tasks, we then analyze the learning behavior of several popular RL methods, i.e., PPO, SAC, and Surprise, on Delivery Mountain Car. Figure 1c shows their learning behavior for the first 2×10^5 time steps. Figure 1c shows that the max height—which the car can reach during an episode regardless of whether the car exhausts the resources or not—is much higher than the max height that the car can reach before exhausting the resources. That is, these methods exhaust the resources when the car only reaches a low height, and then the subsequent exploration of the “unload” action is severely restricted due to the absence of resources. Therefore, these methods have difficulty in achieving the goal of delivering the goods to the mountain top, though they can reach the mountain top. Overall, these methods suffer from inefficient exploration in R3L problems.

6 Methods

In this section, we focus on promoting efficient exploration in R3L problems. We first present a detailed description of our proposed resource-aware exploration bonus (RAEB) and then theoretically analyze the efficiency of RAEB in the finite horizon tabular setting. Due to limited space, we summarize the procedure of RAEB in Appendix B.

6.1 Resource-Aware Exploration Bonus

As shown in Section 5, resource plays an important role in efficient exploration in R3L problems. We observe that the size of accessible state sets of a given state generally positively correlates with its available resources in R3L problems with non-replenishable resources, as states with high resources are inaccessible from states with low resources. Previous empowerment-based exploration methods (Mohamed and Jimenez Rezende 2015; Becker-Ehmck et al. 2021) have shown that exploring states with large accessible state sets is essential for efficient exploration. That is, moving towards states with high resources enables the agent to reach a large number of future states, thus exploring the environment efficiently. However, many existing exploration methods such as Surprise neglect the resource-related information and exhaust resources fast. Thus, they quickly reach states with small accessible state sets and the subsequent exploration is severely restricted.

To promote efficient exploration in R3L problems, we propose a novel resource-aware exploration bonus (RAEB), which encourages the agent to explore novel states with large accessible state sets by promoting resource-saving exploration. Specifically, RAEB takes the form of

$$r_i(s, a) = g(I(s))b(s, a), \quad (2)$$

where $g : \mathbb{R}_+^d \rightarrow \mathbb{R}_+$ is a non-negative function, namely a resource-aware coefficient, and $b(s, a)$ is the exploration bonus to measure the “novelty” of a given state. Intuitively, RAEB encourages the agent to visit unseen states with large accessible state sets. Instantiating RAEB amounts to specifying two design decisions: (1) the measure of “novelty”, (2) and the resource-aware coefficient.

The measure of novelty Surprise-based intrinsic motivation mechanism (Surprise) (Achiam and Sastry 2017) has achieved promising performance in hard exploration high-dimensional continuous control tasks, such as Swimmer-Gather. In this paper, we use surprise—the agent’s surprise about its experiences—to measure the novelty of a given state to encourage the agent to explore rarely seen states. We use the KL-divergence of the true transition probability distribution from a transition model which the agent learns concurrently with the policy as the exploration bonus. Thus, the exploration bonus is of the form

$$b(s, a) = D_{KL}(p \parallel p_\phi)(s, a), \quad (3)$$

where $p : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, \infty)$ is the true transition probability distribution and p_ϕ is the learned transition model. Moreover, researchers (Achiam and Sastry 2017) proposed an approximation to the KL-divergence $D_{KL}(p \parallel p_\phi)(s, a)$,

$$\begin{aligned} D_{KL}(p \parallel p_\phi)(s, a) &= H(p, p_\phi)(s, a) - H(p)(s, a) \\ &\approx H(p, p_\phi)(s, a) \\ &= \mathbb{E}_{s' \sim p(\cdot | s, a)}[-\log p_\phi(s' | s, a)] \end{aligned}$$

The resource-aware coefficient We define $g : \mathbb{R}_+^d \rightarrow \mathbb{R}_+$ in Eqn.(2) as the resource-aware coefficient. We require the function g being an increasing function of the amount of each type of resources, i.e., for each $i \in [d]$, if two resource vector ξ, ξ' satisfy $\xi_i \leq \xi'_i$ and $\xi_j = \xi'_j, j \in [n] \setminus \{i\}$, we have $g(\xi) \leq g(\xi')$. In this paper, we find that a linear function of the quantity of available resources performs well. Suppose the environment has one type of resource, then we define the resource-aware coefficient by

$$g(I(s)) = \frac{I(s) + \alpha}{I_{\max} + \alpha}, \quad (4)$$

where $\alpha > 0$ is a hyperparameter and I_{\max} is the quantity of resources at the beginning of each episode. The hyperparameter α determines the importance of the item $I(s)$ in the resource-aware coefficient and thus controls the degree of resource-saving exploration. Similarly, when the environment involves several types of resources and these resources are relatively independent with each other, we can define the resource-aware coefficient by

$$g(I(s)) = \prod_{i=1}^d \frac{I_i(s) + \alpha_i}{I_{\max, i} + \alpha_i}. \quad (5)$$

Discussion about RAEB We discuss some advantages of RAEB in this part. (1) RAEB makes reasonable usage of resources by promoting novelty-seeking and resource-saving exploration simultaneously. Moreover, RAEB flexibly controls the degree of resource-saving exploration by varying α .

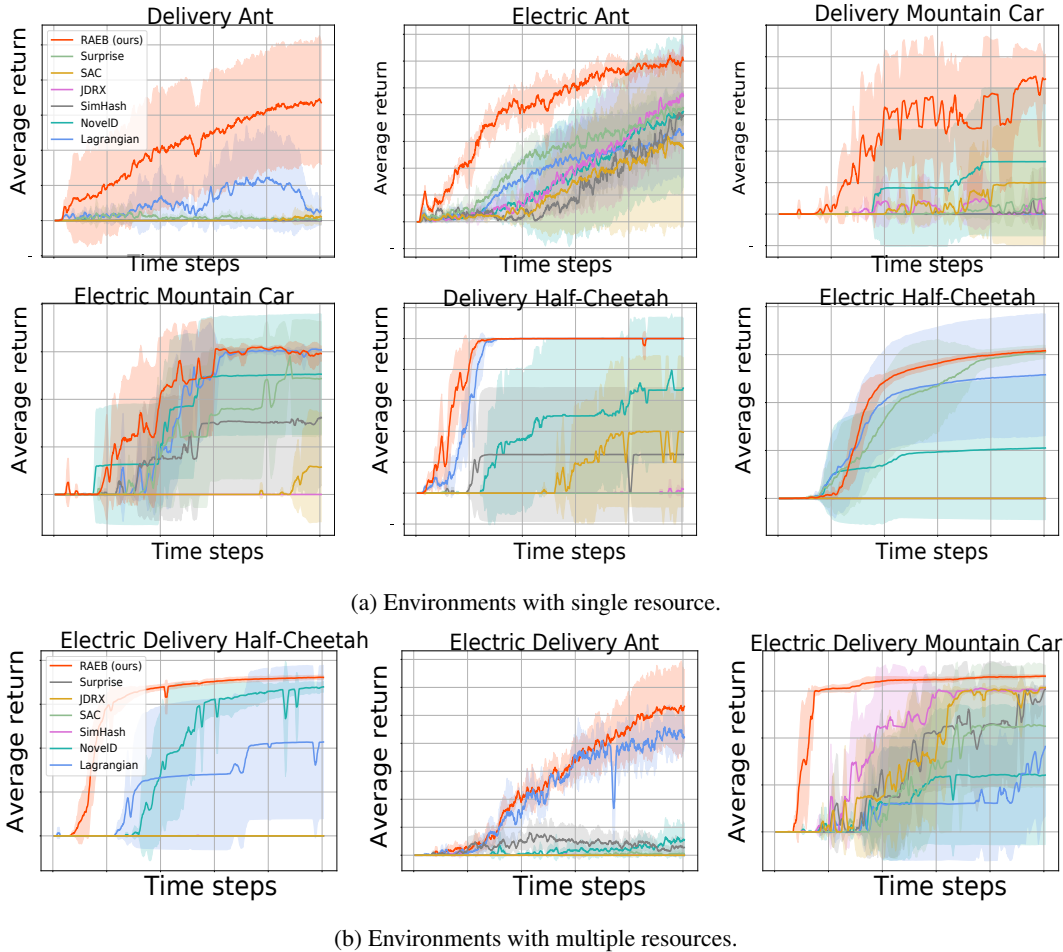


Figure 2: Performance of RAEB and six baselines on nine R3L tasks with limited electricity and/or goods. The solid curves correspond to the mean and the shaded region to the standard deviation over at least five random seeds. For visual clarity, we smooth curves uniformly. The results show that RAEB significantly outperforms these baselines in terms of sample efficiency on several challenging R3L tasks. For Delivery Ant, RAEB improves the sample efficiency by an order of magnitude.

(2) RAEB avoids myopic policies by encouraging the agent to explore novel states with large accessible state sets instead of only exploring novel states.

Theoretical results We provide some theoretical results of RAEB in this part. We restrict the problems to the finite horizon tabular setting and only consider building upon UCB-H (Jin et al. 2018). Researchers (Jin et al. 2018) have shown that Q -learning with UCB-Hoeffding is provably efficient and achieves regret $\tilde{O}(\sqrt{H^4 SAT})$. We view an instantiation of RAEB in the tabular setting as Q -learning with a weighted UCB bonus that multiplies the bounded resource-aware coefficient $\beta(s, a)$ and the UCB exploration bonus. We show that RAEB in the finite horizon tabular setting has a strong theoretical foundation and establishes \sqrt{T} regret.

Theorem 1. *There exists an absolute constant $c > 0$ such that, for any $p \in (0, 1)$, choosing $b_t = c\sqrt{H^3 \iota}/t$, if there exists a positive real d , such that the weights $\beta(s, a) \in [1, d]$ for all state-action pairs (s, a) , then with probability $1 - p$, the total regret of Q -learning with the weighted UCB bonus*

(RAEB) establishes \sqrt{T} regret where $\iota := \log(SAT/p)$.

The details of Q -learning with the weighted bonus algorithm and the rigorous proof to Theorem 1 are in Appendix A. As resources are limited in practice, the resource-aware coefficient defined by Eqn. (4) is bounded. Theorem 1 shows that RAEB in the tabular setting is provably efficient.

7 Experiments

Our experiments have four main goals: (1) Test whether RAEB can significantly outperform state-of-the-art exploration strategies and constrained reinforcement learning methods in R3L tasks. (2) Analyze the effect of each component in RAEB and the sensitivity of RAEB to hyperparameters. (3) Evaluate RAEB variants to provide further insight into RAEB. (4) Illustrate the exploration of RAEB.

As mentioned in Section 5, we design control tasks with limited electricity and/or goods based on Gym (Brockman et al. 2016) and Mujoco (Todorov, Erez, and Tassa 2012). The agent in these tasks can be a continuous mountain car,

a 2D robot (called a ‘half-cheetah’), and a quadruped robot (called an ‘ant’). Then we call these tasks Electric Mountain Car, Electric Half-Cheetah, Electric Ant, Delivery Mountain Car, Delivery Half-Cheetah, Delivery Ant, Electric Delivery Mountain Car, Electric Delivery Half-Cheetah, and Electric Delivery Ant, respectively. For all environments, We augment the states with the available resources. In electric tasks, the agent aims to reach a specific unknown destination with limited electricity. In delivery tasks, the agent aims to deliver the goods to a specific unknown destination. In electric delivery tasks, the agent aims to deliver the goods to a specific unknown destination with limited electricity. Please refer to Appendix B.1 for more details about the environments.

7.1 Evaluation and Comparison Analysis

Baselines. We compare RAEB with several exploration strategies commonly used to solve problems with high-dimensional continuous state-action spaces, such as Ant Maze (Shyam, Jaskowski, and Gomez 2019). The baselines include state-of-the-art (SOTA) exploration strategies and a SOTA constrained reinforcement learning method. For exploration strategies, we compare to Surprise (Achiam and Sastry 2017), a SOTA prediction error-based exploration method, NovelD (Zhang et al. 2021), a SOTA random network based exploration method, Jensen-Rényi Divergence Reactive Exploration (JDRX) (Shyam, Jaskowski, and Gomez 2019), a SOTA information gain based exploration method, SimHash (Tang et al. 2017), a SOTA count-based exploration method, and soft actor critic (SAC) (Haarnoja et al. 2018), a SOTA off-policy algorithm. In terms of constrained reinforcement learning methods, we compare to the Lagrangian method, which uses the consumed resources at each step as a penalty and has achieved strong performance in constrained reinforcement learning tasks (Ray, Achiam, and Amodei 2019). We focus on the model-free reinforcement learning settings in the paper.

Evaluation. We first compare RAEB to the aforementioned baselines in R3L tasks with a single resource, i.e., delivery tasks and tasks with limited electricity. We then compare RAEB to the aforementioned baselines in R3L tasks with multiple resources, i.e., delivery tasks with limited electricity. For all environments, we use the intrinsic reward coefficient $\beta = 0.25$. We use $\alpha = 0.25I_{\max}$ for delivery tasks, $\alpha = 2.5I_{\max}$ for tasks with limited electricity, and $\alpha = [0.25I_{\max}, 2.5I_{\max}]$ for delivery tasks with limited electricity. Note that α controls the degree of resource-saving exploration, and we need to adjust α according to the scarcity of resources in practice.

First, Figure 2a shows that RAEB significantly outperforms the baselines on several challenging R3L tasks with a single resource. For Delivery Ant, we show that the average return of the baselines after 1×10^6 steps is at most 5.65 while RAEB achieves similar performance (11.13) using only 1×10^5 steps. Please refer to Appendix C for detailed results. This result demonstrates that RAEB improves the sample efficiency by an order of magnitude on Delivery Ant. Moreover, some baselines even struggle to attain scores better than random agents on several R3L tasks, which demonstrates that efficient exploration is extremely challenging for




Figure 3: Learning curves on Delivery Ant and Electric Ant. RAEB achieves similar average performance with different β , which shows that RAEB is insensitive to β .




Figure 4: Learning curves on Delivery Ant and Electric Ant. RAEB achieves similar average performance with different α , which demonstrates that RAEB is insensitive to α .

some baselines in R3L tasks. Compared to the Lagrangian method, RAEB still achieves outstanding performance on several challenging R3L tasks. The major reason is that, the constrained reinforcement learning method aims to find policies that satisfy resource constraints, and struggles to make reasonable usage of resources in R3L tasks.

Second, Figure 2b shows that RAEB significantly outperforms the baselines on several R3L tasks with two types of resources. The results highlight the potential of RAEB for efficient exploration in R3L tasks with multiple resources. RAEB scales well in R3L tasks with multiple resources as the two types of resources are independent, and each resource’s resource-aware coefficient encourages the corresponding resource-saving exploration.

7.2 Ablation Study

First, we analyze the sensitivity of RAEB to hyperparameters β and α . Then, We perform ablation studies to understand the contribution of each individual component in RAEB. In this section, we conduct experiments on Delivery Ant and Electric Ant, which have high resolution power due to their difficulty. All results are reported over at least four random seeds. For completeness, we provide additional results in Appendix C.

Sensitivity analysis. First, we compare the performance of RAEB with different reward weighting coefficients β . Intuitively, higher β values should cause more exploration, while too low of an β value reduces RAEB to the base algorithm, i.e., SAC. In the experiments, we set the reward weighting coefficient $\beta = 0.5, 0.25, 0.1$, and 0.05 , respectively. Though choosing a proper reward weighting coefficient remains an open problem for existing intrinsic-reward-based exploration methods (Burda et al. 2019), the results in Figure




Figure 5: Learning curves on Delivery Ant and Electric Ant. The results show that each component of RAEB is significant for efficient exploration in R3L tasks.




Figure 6: We compare RAEB to two RAEB variants on Delivery Ant and Electric Ant. The results show that these variants struggle to achieve outstanding performance in both environments compared to RAEB.

3 show that there is a wide β range for which RAEB achieves comparable average performance on Delivery Ant and Electric Ant. Then, we compare the performance of RAEB with different α on Delivery Ant and Electric Ant. The results in Figure 4 show that there is a wide α range for which RAEB achieves competitive performance on Delivery Ant and Electric Ant. Overall, the results in Figures 3 and 4 show that RAEB is insensitive to hyperparameters β and α .

Contribution of each component. RAEB contains two components: the measure of novelty and the resource-aware coefficient. We perform ablation studies to understand the contribution of each individual component. Figure 5 shows that the full algorithm RAEB significantly outperforms each ablation with only a single component on Delivery Ant and Electric Ant. On delivery Ant, RAEB without resources and RAEB without surprise even struggle to perform better than a random agent. The results show that each component is significant for efficient exploration in these R3L tasks. Though the surprise bonus encourages novelty-seeking exploration, it tends to exhaust resources fast. And only the resource-aware coefficient tends to encourage the agent to save resources without consuming resources, while resource-consuming exploration is significant in R3L tasks.

7.3 Compared to RAEB Variants

Based on the idea of preventing the agent from exhausting resources fast, it is natural to use the quantity of the remaining resources as a reward bonus instead of a resource-aware coefficient. We use the bonus to train SAC and Surprise, called SAC with resources bonus (SACRB) and Surprise with resources bonus (SurpriseRB), respectively. Specifically, SACRB and SurpriseRB add an additional reward

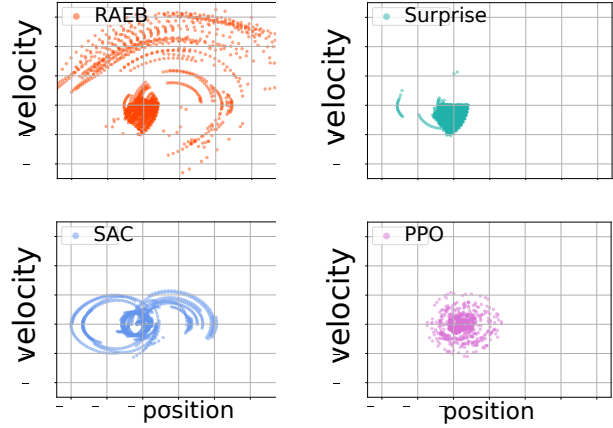


Figure 7: Illustration of RAEB exploration on Delivery Mountain Car. Colored points illustrate the states (position and velocity of the car) where the agent consumes resources in 0.2 million steps.

bonus $c \frac{I(s)+\alpha}{I_{\max}+\alpha}$ to the reward. We compare RAEB to these variants on Delivery Ant and Electric Ant. Figure 6 demonstrates that RAEB significantly outperforms these variants in both two environments, which shows the superiority of RAEB compared to its variants. For Delivery Ant, the RAEB variants even struggle to perform better than random agents. The major advantage of RAEB over SurpriseRB is that it is able to well combine the exploration ability of the surprise bonus and the resource-aware bonus by using a resource-aware coefficient instead of adding an additional bonus, as the scale between the surprise and the resource-aware bonus can be extremely different. Moreover, SurpriseRB is sensitive to the hyperparameter c . (See Appendix C)

7.4 Illustration of RAEB Exploration

To illustrate the exploration in the Delivery Mountain Car environment, we visualize the states where the agent unloads the goods in Figure 7. Notice that we visualize the position and velocity of the car while ignoring the quantity of remaining resources for conciseness. Figure 7 shows that the agents trained by PPO, SAC, and Surprise unload the goods in the same states repeatedly, leading to inefficient exploration. Moreover, Figure 7 shows that RAEB significantly reduces unnecessary resource-consuming trials, while effectively encouraging the agent to explore unvisited states.

8 Conclusion

In this paper, we first formalize the decision-making tasks with resources as a resource-restricted reinforcement learning, and then propose a novel resource-aware exploration bonus to make reasonable usage of resources. Specifically, we quantify the RAEB of a given state by both the measure of novelty and the quantity of available resources. We conduct extensive experiments to demonstrate that the proposed RAEB significantly outperforms state-of-the-art exploration strategies on several challenging robotic delivery and autonomous electric robot tasks.

Acknowledgments

We would like to thank all the anonymous reviewers for their insightful comments. This work was supported in part by National Science Foundations of China grants U19B2026, U19B2044, 61836006, and 62021001, and the Fundamental Research Funds for the Central Universities grant WK3490000004.

References

- Achiam, J.; Held, D.; Tamar, A.; and Abbeel, P. 2017. Constrained Policy Optimization. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, 22–31.
- Achiam, J.; and Sastry, S. 2017. Surprise-Based Intrinsic Motivation for Deep Reinforcement Learning. *CoRR*, abs/1703.01732.
- Becker-Ehmck, P.; Karl, M.; Peters, J.; and van der Smagt, P. 2021. Exploration via Empowerment Gain: Combining Novelty, Surprise and Learning Progress. In *ICML 2021 Workshop on Unsupervised Reinforcement Learning*.
- Bellemare, M.; Srinivasan, S.; Ostrovski, G.; Schaul, T.; Saxton, D.; and Munos, R. 2016. Unifying Count-Based Exploration and Intrinsic Motivation. In *Advances in Neural Information Processing Systems 29*.
- Bhatia, A.; Varakantham, P.; and Kumar, A. 2019. Resource constrained deep reinforcement learning. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 29, 610–620.
- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. Openai gym. *arXiv preprint arXiv:1606.01540*.
- Burda, Y.; Edwards, H.; Pathak, D.; Storkey, A.; Darrell, T.; and Efros, A. A. 2019. Large-Scale Study of Curiosity-Driven Learning. In *International Conference on Learning Representations*.
- Chen, H.; Lam, H.; Li, F.; and Meisami, A. 2020. Constrained Reinforcement Learning via Policy Splitting. In *Proceedings of The 12th Asian Conference on Machine Learning*, volume 129 of *Proceedings of Machine Learning Research*, 209–224.
- Chung, J.; Lawrance, N.; and Sukkarieh, S. 2014. Learning to soar: Resource-constrained exploration in reinforcement learning. *The International Journal of Robotics Research*, 34: 158–172.
- Cui, D.; Peng, Z.; Xiong, J.; Xu, B.; and Lin, W. 2020. A Reinforcement Learning-Based Mixed Job Scheduler Scheme for Grid or IaaS Cloud. *IEEE Transactions on Cloud Computing*, 8(4): 1030–1039.
- Ding, D.; Zhang, K.; Başar, T.; and Jovanović, M. 2017. Natural Policy Gradient Primal-Dual Method for Constrained Markov Decision Processes. In *Advances in Neural Information Processing Systems 30*, 2753–2762.
- Grizzle, J. W.; Chevallereau, C.; Sinnenet, R. W.; and Ames, A. D. 2014. Models, feedback control, and open problems of 3D bipedal robotic walking. *Automatica*, 50(8): 1955–1988.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. volume 80, 1861–1870.
- Houthoofd, R.; Chen, X.; Chen, X.; Duan, Y.; Schulman, J.; De Turck, F.; and Abbeel, P. 2016. VIME: Variational Information Maximizing Exploration. In Lee, D.; Sugiyama, M.; Luxburg, U.; Guyon, I.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 29, 1109–1117.
- Jin, C.; Allen-Zhu, Z.; Bubeck, S.; and Jordan, M. I. 2018. Is Q-Learning Provably Efficient? In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 31*, 4863–4873.
- Kempka, M.; Wydmuch, M.; Runc, G.; Toczek, J.; and Jaśkowski, W. 2016. Vizdoom: A doom-based ai research platform for visual reinforcement learning. In *2016 IEEE Conference on Computational Intelligence and Games (CIG)*, 1–8. IEEE.
- Kormushev, P.; Calinon, S.; and Caldwell, D. G. 2013. Reinforcement learning in robotics: Applications and real-world challenges. *Robotics*, 2(3): 122–148.
- Kormushev, P.; Ugurlu, B.; Calinon, S.; Tsagarakis, N. G.; and Caldwell, D. G. 2011. Bipedal walking energy minimization by reinforcement learning with evolving policy parameterization. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 318–324. IEEE.
- Mohamed, S.; and Jimenez Rezende, D. 2015. Variational Information Maximisation for Intrinsically Motivated Reinforcement Learning. In Cortes, C.; Lawrence, N.; Lee, D.; Sugiyama, M.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.
- Pathak, D.; Agrawal, P.; Efros, A. A.; and Darrell, T. 2017. Curiosity-driven Exploration by Self-supervised Prediction. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, 2778–2787.
- Ray, A.; Achiam, J.; and Amodei, D. 2019. Benchmarking safe exploration in deep reinforcement learning. *arXiv preprint arXiv:1910.01708*, 7(1): 2.
- Resnick, C.; Eldridge, W.; Ha, D.; Britz, D.; Foerster, J.; Togelius, J.; Cho, K.; and Bruna, J. 2018. Pommerman: A multi-agent playground. *arXiv preprint arXiv:1809.07124*.
- Schulman, J.; Levine, S.; Abbeel, P.; Jordan, M.; and Moritz, P. 2015. Trust region policy optimization. In *International conference on machine learning*, 1889–1897. PMLR.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. abs/1707.06347.
- Shyam, P.; Jaskowski, W.; and Gomez, F. 2019. Model-Based Active Exploration. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, 5779–5788. PMLR.

Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; Bolton, A.; et al. 2017. Mastering the game of go without human knowledge. *nature*, 550(7676): 354–359.

Strehl, A. L.; and Littman, M. L. 2004. An empirical evaluation of interval estimation for Markov decision processes. In *16th IEEE International Conference on Tools with Artificial Intelligence*, 128–135.

Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement learning: An introduction*. MIT press.

Tang, H.; Houthoofd, R.; Foote, D.; Stooke, A.; Xi Chen, O.; Duan, Y.; Schulman, J.; DeTurck, F.; and Abbeel, P. 2017. #Exploration: A Study of Count-Based Exploration for Deep Reinforcement Learning. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

Tesauro, G.; et al. 2005. Online resource allocation using decompositional reinforcement learning. In *AAAI*, volume 5, 886–891.

Todorov, E.; Erez, T.; and Tassa, Y. 2012. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 5026–5033.

Vinyals, O.; Babuschkin, I.; Czarnecki, W. M.; Mathieu, M.; Dudzik, A.; Chung, J.; Choi, D. H.; Powell, R.; Ewalds, T.; Georgiev, P.; et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782): 350–354.

Zhang, T.; Xu, H.; Wang, X.; Wu, Y.; Keutzer, K.; Gonzalez, J. E.; and Tian, Y. 2021. NovelD: A Simple yet Effective Exploration Criterion. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 25217–25230. Curran Associates, Inc.

Zhang, W.; and Dietterich, T. G. 1995. A reinforcement learning approach to job-shop scheduling. In *IJCAI*, volume 95, 1114–1120. Citeseer.