

UC Irvine

UC Irvine Electronic Theses and Dissertations

Title

Curiosity-Driven Multi-Criteria Hindsight Experience Replay

Permalink

<https://escholarship.org/uc/item/3pg2h1b3>

Author

Lanier, John Banister

Publication Date

2019

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA,
IRVINE

Curiosity-Driven Multi-Criteria Hindsight Experience Replay

THESIS

submitted in partial satisfaction of the requirements
for the degree of

MASTER OF SCIENCE

in Computer Science

by

John Banister Lanier

Thesis Committee:
Professor Pierre Baldi, Chair
Professor Sameer Singh
Professor Eric Mjolsness

2019

TABLE OF CONTENTS

	Page
LIST OF FIGURES	iv
LIST OF TABLES	v
ACKNOWLEDGMENTS	vi
ABSTRACT OF THE THESIS	vii
1 Introduction	1
1.1 Related Work	3
1.1.1 Curiosity-Driven Exploration	3
1.1.2 Curriculum Learning in Goal-Based Tasks	4
1.1.3 Hindsight methods	5
1.1.4 Block Stacking	5
1.2 Background	6
1.2.1 Reinforcement Learning	6
1.2.2 DDPG	7
1.2.3 DDPG with Goals	8
1.2.4 Hindsight Experience Replay	8
1.3 Environments	9
2 Methods	11
2.1 Curiosity Driven Exploration with Multiple Policies	12
2.1.1 Exploit Actor and Critic	14
2.1.2 Explore Actor and Critic	14
2.1.3 Combined Actor and POP-ART	15
2.2 Multi-Criteria Hindsight	18
2.3 Curriculum	19
3 Experiments	21
4 Conclusion	27
Bibliography	28

Appendices	32
A Links	32
B Experiment Details	33

LIST OF FIGURES

	Page
1.1 Simulated Robotic Block Stacking Environment	2
2.1 Forward Network Connections for DDPG+HER Learner with Curiosity-Driven Exploration	13
2.2 DDPG+HER with Curiosity Algorithm	16
2.3 Multi-Criteria HER Augmentation Step Algorithm	19
3.1 Success Rates and Per-Episode Rewards for Block Stacking with Binary Rewards	23
3.2 Success Rates and Per-Episode Rewards for Block Stacking with Incremental Rewards	23

LIST OF TABLES

	Page
3.1 Highest Success Rates with Binary Rewards over 100 Episode Sliding Window	24
3.2 Highest Success Rates with Incremental Rewards over 100 Episode Sliding Window	24

ACKNOWLEDGMENTS

I would like to thank Stephen McAleer for his assistance and guidance with regard to this work.

I would also like to thank Alexander Shmakov and Forest Agostinelli for their advice and interesting conversations on reinforcement learning.

Finally, I would like to thank Professor Pierre Baldi for facilitating and allowing me to participate in a constructive lab environment, making this work possible.

ABSTRACT OF THE THESIS

Curiosity-Driven Multi-Criteria Hindsight Experience Replay

By

John Banister Lanier

Master of Science in Computer Science

University of California, Irvine, 2019

Professor Pierre Baldi, Chair

Dealing with sparse rewards is a longstanding challenge in reinforcement learning. The recent use of hindsight methods have achieved success on a variety of sparse-reward tasks, but they fail on complex tasks such as stacking multiple blocks with a robot arm in simulation. Curiosity-driven exploration using the prediction error of a learned dynamics model as an intrinsic reward has been shown to be effective for exploring a number of sparse-reward environments. We present a method that combines hindsight with curiosity-driven exploration and curriculum learning in order to solve the challenging sparse-reward block stacking task. We are the first to stack more than two blocks using only sparse reward without human demonstrations.

Chapter 1

Introduction

Goal-based reinforcement learning has become an important framework for formulating and solving goal-based sequential decision making tasks. In goal-based reinforcement learning, the agent’s rewards are usually dependent on achieving a goal, and it chooses its actions using a goal-conditioned policy. Goal-conditioned policies can enable a reinforcement learning agent to generalize to new goals after training on a many different goals in the same environment [36].

Goal-based reinforcement learning environments can be given a binary and sparse reward that is encountered only when the goal is reached. Defining reward in this way ensures that if the agent maximizes reward then it also reaches the user’s intended goal, which is not necessarily true of manually-shaped dense rewards [41]. However, sparse rewards are also difficult to learn from. As the length of a sparse-reward task increases, it becomes less likely that an agent will discover how to reach its goal through random exploration [37]. This problem is exacerbated when a sparse reward depends on the fulfillment of multiple goals or criteria.

Recently, hindsight methods have served as a popular solution to sparse-reward goal-oriented learning by training an agent on the goals that it actually reached in addition to those

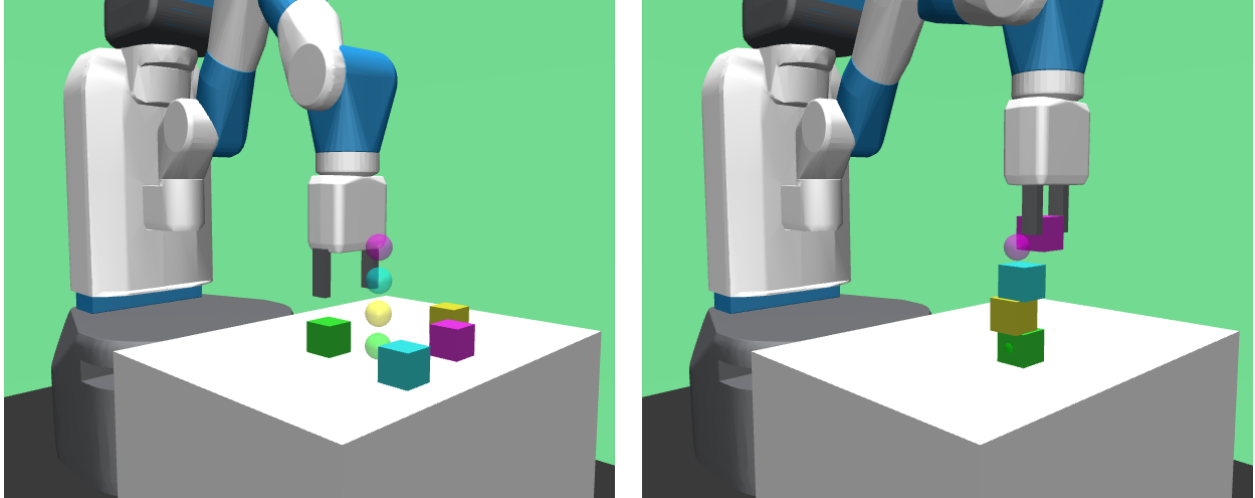


Figure 1.1: The simulated robotic block stacking environment. A goal consists of target positions where the blocks need to be placed, shown here as colored spheres. In the incremental reward environment, the agent receives a reward for each block being in its target position. In the sparse reward environment, the agent receives a reward only when every block is in its target position.

which were intended [1]. This is done in the hope that knowledge of how to reach randomly discovered goals will allow an agent to generalize well enough to find its assigned goals. However, in many environments, an agent can be asked to reach goals that are very different from those it may discover by chance, causing such generalization to be difficult. In these cases, the same sparse reward issues remain, making it challenging for an agent to learn how to accomplish its given objectives.

Stacking multiple blocks in a simulated robotics environment is a sparse-reward, goal-based task that highlights shortcomings of hindsight learning. Multiple-block stacking is too difficult for established hindsight methods like Deep Deterministic Policy Gradients with Hindsight Experience Replay (DDPG+HER) to reliably solve without access to human demonstrations [29]. Satisfying all criteria of a block stacking goal requires learning multiple skills to correctly place each block, and the end goals are very different from those that the agent may discover with random exploration. Achieving reward by correctly placing all blocks is precarious and requires long chains of specific actions. Therefore, under a sparse reward, even with hindsight, it is highly unlikely that an agent will discover the complex sequences of actions required to

place every block in its correct position on the stack. Our method is the first method that is able to solve sparse-reward block stacking for more than two blocks without access to human demonstrations.

To solve sparse-reward multi-block stacking without help from demonstration, we use DDPG+HER combined with curiosity-driven exploration and curriculum learning. In order to balance improved exploration with exploitation during training, we introduce a new method of combining data from both curiosity-based and standard policies in an off-policy fashion. Additionally, we introduce a form of hindsight experience replay that is more sample efficient for multi-criteria goal-based environments. We show that the advantages introduced by each of these methods complement the others, and that the combination of all of them is necessary to solve the hardest stacking tasks.

1.1 Related Work

1.1.1 Curiosity-Driven Exploration

We refer to curiosity-driven exploration as any method that attempts to drive an agent to explore trajectories which it has not visited frequently before, usually by making the agent pursue some form of exploration related objective or reward.

Curiosity-driven exploration has been approached by training agents to maximize information gain [26, 20], pursue less visited areas using state pseudo-counts [2, 30], and maximize state empowerment [18, 28, 23].

We focus on exploration by performing actions that both challenge and improve an agent’s ability to model the world [21, 17, 38, 39]. We approach this by training a dynamics model on the state transitions that our agent visits and encouraging the agent to maximize the model’s

per-sample error on those transitions. Assuming a dynamics model is more accurate on transitions that it has seen frequently before, such an agent seeking to challenge the dynamics model should be inclined to visit new, rarely before seen state transitions. Choosing actions to directly challenge an online trained dynamics model has been shown to result in complex emergent behaviors [19]. Using a dynamics model’s error as an RL exploration reward can motivate an agent to seek out novel states, sometimes solving an environment’s objective without extrinsic rewards, and combining environmental rewards with a bonus exploration reward has the potential to increase an agent’s learning speed and end-performance [31, 4]. On the same note, training a model to predict the output of a random function from state features and choosing actions to maximize its error helped achieve state-of-the art performance on the Montezuma’s Revenge Atari domain [5].

1.1.2 Curriculum Learning in Goal-Based Tasks

Previous applications of curriculum learning [3] to goal based environments include training on a variety of tolerances for considering goals achieved [16], masking certain goal dimensions to allow all such values on an axis to be sufficient for success [13], and generating curricula that walk backwards from a predefined success state [14, 27].

Intrinsically motivated goal exploration processes (IMGEPs) have also been used to automatically generate goals which maximize learning progress across one [15, 32, 24] or multiple [6] tasks.

1.1.3 Hindsight methods

Our work builds on Hindsight Experience Replay (HER) [1] as a way to effectively augment goal oriented transition samples for a replay buffer. Hindsight has also been adapted to policy gradient settings [36].

Efforts have been made to increase the efficiency of HER by prioritizing the sampling of more relevant transitions. This has been done by attributing higher importance to transitions and trajectories in which more physical work is done by the agent [43], rare goal states are achieved [42], or higher temporal difference error is measured [9].

1.1.4 Block Stacking

Stacking multiple blocks with sparse rewards has been solved before using expert demonstration in [11] and [29]. Our work is a direct followup to the latter, as we solve a similar set of environments without demonstration.

Stacking only 2-blocks with sparse rewards has been solved without demonstration by training on an automatic curriculum which selects tasks from a small collection, prioritizing tasks with higher changes in learning progress [6] and by collecting data from multiple policies following auxiliary objectives to accomplish predefined interesting actions [37]. Dense reward robotic block stacking tasks have been solved before using both a model-based approach, PILCO, [8, 7] and by initializing the environment at intermediate stages of the task [35].

1.2 Background

1.2.1 Reinforcement Learning

We consider the standard reinforcement learning formalism in which an agent interacts with an environment E . The environment is fully observable, and consists of a set of states S , a set of actions A , a reward function $r : S \times A \rightarrow \mathbb{R}$, an initial state distribution $p(s_0)$ and transition dynamics $p(s_{t+1}|s_t, a_t)$. At each timestep t , the agent observes a state s_t , takes an action a_t , and receives a reward r_t . The agent chooses these actions using a policy π , which is a conditional distribution over actions given states. In this paper, we consider deterministic policies which map directly from states to actions $\pi : S \rightarrow A$.

The discounted sum of future rewards is defined as the return $R_t = \sum_{i=t}^T \gamma^{(i-t)} r_i$ over some time horizon T and with a discounting factor $\gamma \in [0, 1]$. We define ρ^π as the state visitation distribution when taking actions according to π . The goal in reinforcement learning is to learn a policy π to maximize the expected return $J = \mathbb{E}_{s_i \sim \rho^\pi, a_i \sim \pi, r_i \sim E} [R_0 | s_0]$.

The expected return when taking actions according to a specific policy π is called the Q-function or action-value function, and is defined as:

$$Q^\pi(s_t, a_t) = \mathbb{E}_{s_{i>t} \sim \rho^\pi, r_{i \geq t} \sim E} [R_t | s_t, a_t] \tag{1.1}$$

which can be recursively stated as the Bellman equation:

$$Q^\pi(s_t, a_t) = \mathbb{E}_{r_t, s_{t+1} \sim E} [r_t + \gamma Q^\pi(s_{t+1}, \pi(s_t, a_t))] \tag{1.2}$$

Because π is deterministic, the expectation in equation 1.2 depends only on the environment, allowing off-policy methods to learn Q^π while using transitions generated with some other stochastic policy β .

1.2.2 DDPG

Our work uses the Deep Deterministic Policy Gradients algorithm (DDPG) [25], which is an off-policy, model-free reinforcement learning algorithm designed for use with deep neural networks in continuous action spaces. DDPG uses an actor-critic methodology. Two neural networks are trained: a critic $Q : S \times A \rightarrow \mathbb{R}$ parameterized by θ^Q , and an actor serving as the policy $\pi : S \rightarrow A$, which is updated using the policy gradient to directly maximize Q^π with respect to the policy’s parameters θ^π :

$$\nabla_{\theta^\pi} J = \mathbb{E}_{s_t \sim \rho^\beta} [\nabla_{\theta^\pi} Q(s_t, a | \theta^Q) |_{a=\pi(s_t | \theta^\pi)}] \quad (1.3)$$

This quantity can be estimated with the following:

$$\nabla_{\theta^\pi} J \approx \frac{1}{N} \sum_i \nabla_a Q(s_i, a | \theta^Q) |_{a=\pi(s_i)} \nabla_{\theta^\pi} \pi(s_i | \theta^\pi) \quad (1.4)$$

The critic’s parameters θ^Q are updated to minimize the loss:

$$L^{crit} = \frac{1}{N} \sum_i (y_i - Q(s_i, a_i | \theta^Q))^2 \quad (1.5)$$

where

$$y_t = r_t + \gamma Q'(s_{t+1}, \pi'(s_{t+1})) \quad (1.6)$$

For stability, slower moving target networks π' and Q' are used to calculate y_t . These network’s parameters are exponential moving averages of θ^π and θ^Q respectively.

DDPG maintains a replay buffer R containing transition samples, which are tuples (s_t, a_t, r_t, s_{t+1}) , and alternates between two stages. The first stage is to gather experience for R by performing rollouts on the environment, choosing actions from a new policy $\beta = \pi + \epsilon$ where ϵ is random.

The second stage is to train π and Q on batches of transition samples from R .

To efficiently gather experience, we run DDPG in parallel using multiple workers with synchronized copies of each network, averaging parameters across workers after each update.

1.2.3 DDPG with Goals

In our work, we follow a goal based-framework. A goal $g \in G$ is sampled each episode, and π and Q are conditioned on these goals, making them $\pi : S \times G \rightarrow A$ and $Q : S \times A \times G \rightarrow \mathbb{R}$. Furthermore, the replay buffer instead stores transition samples as $(s_t || g, a_t, r_t, s_{t+1} || g)$, where the states are each concatenated with a goal. The environments' reward functions $r_t = r_{env}(s_{t+1}, g)$ are also parameterized on whether a new state meets these goals.

1.2.4 Hindsight Experience Replay

In goal-based scenarios, hindsight experience replay increases the sample efficiency of replay buffer based algorithms like DDPG by adding additional augmented samples to the replay buffer. In doing so, HER allows the agent to evaluate its progress not only towards the goals that it was given by the environment, but also towards those that it actually reached in experience gathering rollouts, thus giving the agent hindsight.

HER acts by duplicating transition samples before placing them in the replay buffer, and in those duplicates, augmenting them by replacing the environment-provided goals with goals that were actually reached later in the same episode. HER requires the learning algorithm to have access to the reward function, and the rewards in the augmented samples are updated according to the newly replaced goals.

HER can also be implemented by expressly storing unmodified transition samples in the

replay buffer and, with a certain probability, augmenting them when they are sampled from it. We use this method in our work.

1.3 Environments

The block stacking environments that we consider in this work are based on the Fetch robot environments from the OpenAI Gym API [33] and are similar those used in [29]. We test on separate environments for $n = 2$ to 4 blocks. Each episode, target block locations are initialized in a stack somewhere on the surface of a table. The n blocks are initialized at random locations on the table away from the target stack location. The blocks are uniquely labeled, and each block always goes to the same vertical position in the stack. The agent has $25n$ timesteps before the environment resets. These environments are fully observable, and observations include the claw state and full position, rotation, and velocity for both the robot’s gripper and each block. These environments’ goals specify the target positions of each block, and a block is considered correctly placed if its position is within an error tolerance e from its target position. Actions are continuous and control the robot gripper’s movement in 3 dimensions as well as the state of the claw.

Similar to [29], we consider two sparse reward formulations with these environments: binary and incremental. We can provide a single binary reward when the goal is fully achieved upon correctly placing all blocks:

$$r_t^{binary} = \begin{cases} 0 & \text{all blocks in place} \\ -1 & \text{otherwise} \end{cases} \quad (1.7)$$

We also consider incremental rewards for each block correctly placed:

$$r_t^{incremental} = \text{no. of blocks in place} - \text{no. of blocks} \quad (1.8)$$

In both cases, we also add 1 to the reward for moving the gripper away from the blocks once they are all correctly placed, but this does not affect whether any goal is considered achieved.

Chapter 2

Methods

To solve multi-block stacking with both incremental and binary rewards, we use three methods to improve the performance of a standard multi-worker DDPG+HER learner: curiosity-driven exploration, multi-criteria HER, and curriculum learning.

First, we incorporate curiosity-driven exploration by training a forward dynamics model on state transitions visited by the agent and treating the dynamics model’s prediction error on these transitions as an exploration reward. In order to have a certain portion of workers explore while others exploit, we train three separate policies to maximize exploration rewards, environmental rewards, and a weighted combination of both. Experience from rollouts is shared among each network regardless of which policy collected it. By doing so, we can use different policies at training time than at test time.

Second, we introduce a form of hindsight experience replay better suited for multi-criteria goal-based environments, where a criteria in our environment is defined as the position of a specific block. Our method randomly performs the goal replacement operation on each independent criteria in a goal rather than on an entire goal at once, decoupling the individual effects of each criteria on the reward function and providing higher sample efficiency.

Third, we use curriculum learning by training the agent on two easier skill-building environments before training on the target multi-block stacking task.

2.1 Curiosity Driven Exploration with Multiple Policies

We use curiosity-driven exploration to encourage an agent to visit transitions which are novel and surprising to it. We define an auxiliary exploration reward in addition to environmental reward, and we train separate critics for each. The explore critic Q_e predicts the action-value function for exploration rewards, and the exploit critic Q_r predicts the action-value function for environmental rewards. We train three actor policies π_e, π_r, π_c which respectively maximize exploration rewards, environmental rewards, and a weighted combination of both. By training separate policies, we can make our agent pursue multiple and various objectives at training time and maximize only environmental rewards at test time.

We maintain a forward dynamics neural network $D : S \times A \rightarrow S$ parameterized by θ^d to predict the next observation given the current observation and action, and we train it on the same transition samples from the replay buffer as our agent at each DDPG update step. For each transition sample trained on, an exploration reward for the sample is defined as the squared error between the predicted next state $D(s_t, a_t | \theta^d)$ and the actual next state s_{t+1} . The minibatch loss function for D and the exploration reward is formulated as:

$$L_d = \frac{1}{N} \sum_i r_i^{explore} = \frac{1}{N} \sum_i (s_{i+1} - D(s_i, a_i | \theta^d))^2 \quad (2.1)$$

By passing this error $r_i^{explore}$ to our agent as an exploration reward to maximize, we encourage our agent to pursue transitions that are difficult to predict and unlike transitions currently in the replay buffer.

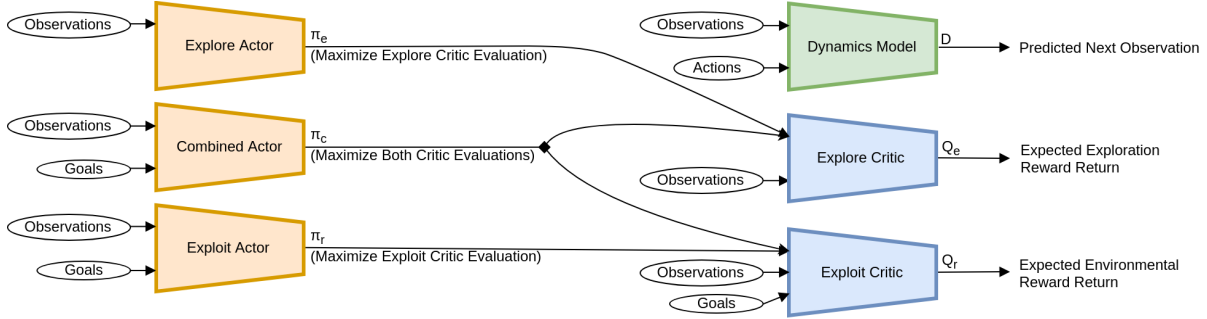


Figure 2.1: Forward Network Connections for DDPG+HER Learner with Curiosity-Driven Exploration. During testing, actions are taken by the exploit actor to maximize return on environmental rewards. During training, actions can be taken by any actor depending on which objectives we wish to emphasize.

With two separate reward sources, we group our multiple actors and critics into two DDPG actor-critic pairs. On exploration reward, we train Q_e and π_e as our explore actor-critic pair. On environmental reward, we train Q_r and π_r as our exploit actor-critic pair. In addition to these actor-critic pairs, we also train π_c as our combined actor. π_c pursues both exploration and environmental reward by maximizing a weighted average of both critics’ action-value functions.

We train multiple actors towards different objectives so that we can assign a portion of our workers to follow an exploration related policy π_e or π_c while the rest follow the exploit policy π_r . Doing so allows us to diversify the experience gathered and make less sacrifices toward either exploration or exploitation objectives than if we were to only ever choose actions which maximize a weighted combination of the two. With multiple actors, we can specialize our workers and maximize both environmental and exploration rewards when gathering experience and then use π_r at test time to solely maximize environmental rewards. Below we describe these networks in more detail.

2.1.1 Exploit Actor and Critic

$\pi_r : S \times G \rightarrow A$ and $Q_r : S \times G \times A \rightarrow \mathbb{R}$ together form the exploit actor-critic pair, which is trained on the normal DDPG+HER goal-based RL objective for maximizing return on environmental rewards conditioned on goals. This actor-critic pair follows the same configuration and update rules as what would be used in vanilla DDPG+HER. The loss function for Q_r to minimize with respect to its parameters θ_r^Q is:

$$L_r^{crit} = \mathbb{E}_{s_t \sim \rho^\beta, a_t \sim \beta, r_t \sim E} (y_t^r - Q_r(s_t, a_t, g | \theta_r^Q))^2 \quad (2.2)$$

where y_t^r is calculated using the target exploit actor and critic π_r' and Q_r' :

$$y_t^r = r_t^{env} + \gamma Q_r'(s_{t+1}, \pi_r'(s_{t+1}, g), g) \quad (2.3)$$

π_r is updated using the standard goal-based DDPG policy gradient to maximize Q_r with respect to π_r 's parameters θ_r^π .

2.1.2 Explore Actor and Critic

$\pi_e : S \rightarrow A$ and $Q_e : S \times A \rightarrow \mathbb{R}$ together form the explore actor-critic pair, which is trained on the objective of maximizing return on exploration reward. Goals do not affect exploration rewards and are not factored in these calculations. The loss function for Q_e to minimize with respect to its parameters θ_e^Q is:

$$L_e^{crit} = \mathbb{E}_{s_t \sim \rho^\beta, a_t \sim \beta, r_t \sim E} (y_t^e - Q_e(s_t, a_t | \theta_e^Q))^2 \quad (2.4)$$

where y_t^e is calculated using the target explore actor and critic π_e' and Q_e' :

$$y_t^e = r_t^{explore} + \gamma Q_e'(s_{t+1}, \pi_e'(s_{t+1})) \quad (2.5)$$

Likewise, π_e is updated using the standard DDPG policy gradient to maximize Q_e with respect to π_e 's parameters θ_e^π .

2.1.3 Combined Actor and POP-ART

Once our agent has an idea of how to find environmental rewards, it is usually more advantageous to explore trajectories close to what actually results in those rewards. Towards this end, we train our combined actor $\pi_c : S \times G \rightarrow A$ to choose actions that maximize both the exploration and exploitation objectives simultaneously. π_c outputs actions that maximize the weighted combination of both Q_e and Q_r 's action-value functions.

We intend to maintain a normalized scale at which to compare the return estimates from Q_e and Q_r so that we can intuitively weight their relative importance to π_c . We also need to account for the fact that the magnitude of both action-value functions may change drastically over the course of training. This is especially true of Q_e which predicts the return from the moving exploration reward function. To accomplish both of these goals, each of the targets y^e and y^r for Q_e and Q_r are adaptively normalized such that we also maintain normalized versions n_e^Q and n_r^Q of both action-value functions with the same relative scale at all times. We can then intuitively weight the relative importance of n_e^Q and n_r^Q for π_c to maximize. In our case, we weight them equally.

To do this, we use PopArt normalization [40], which allows us to adaptively normalize our critics' targets without hurting the accuracy of our predictions. Here we only sketch PopArt informally. See [40] for more details. For each critic target y^e and y^r we keep an online

Algorithm 1: DDPG+HER with Curiosity

Given: Worker policies $\pi_0, \pi_1, \dots, \pi_w | \pi_i \in \{\pi_e, \pi_r, \pi_c\}$
Randomly initialize networks $D, n_e^Q, n_r^Q, \pi_e, \pi_r, \pi_c$
Initialize target networks $n_e^{Q'}, n_r^{Q'}, \pi_e', \pi_r', \pi_c'$

(Execute for each parallel worker i):

Initialize replay buffer R

for $Epoch = 1, \dots, E$ **do**

for $Cycle = 1, C$ **do**

for $Episode = 1, M$ **do**

 Sample ϵ and set $\beta \leftarrow \pi_i + \epsilon$

 Receive initial state s_0 and goal g

for $t = 0, T$ **do**

 Select action $a_t = \beta(s_t, g)$ with noise

 Take action a_t , receive r_t, s_{t+1}

 Store $(s_t || g, a_t, r_t, s_{t+1} || g)$ in R

end

end

for $Batch = 1, \dots, K$ **do**

 Sample batch B from R with HER augmentations

 Train D on B

foreach transition sample j in B **do**

 Set $r_j^{explore}$ and add it to sample

end

 Train $n_e^Q, n_r^Q, \pi_e, \pi_r, \pi_c$ on B

 Update target networks

 Average network parameters over workers

end

end

 Test performance on episodes using π_r

end

Figure 2.2: DDPG+HER with Curiosity Algorithm

estimate of its mean and standard deviation σ_e, μ_e and σ_r, μ_r . We then parameterize Q_e and Q_r as linear transformations of the suitably normalized action-value functions n_e^Q and n_r^Q :

$$\begin{aligned} Q_e(s, a | \theta_e^Q) &= \sigma_e n_e^Q(s, a | \theta_e^Q) + \mu_e \\ Q_r(s, a, g | \theta_r^Q) &= \sigma_r n_r^Q(s, a, g | \theta_r^Q) + \mu_r \end{aligned} \tag{2.6}$$

n_e^Q and n_r^Q are the actual networks that we train parameterized by θ_e^Q and θ_r^Q , and when the statistics σ_e, μ_e and σ_r, μ_r are updated, the top layers of n_e^Q and n_r^Q are also adjusted to preserve equation 2.6. Similarly, our target critics Q'_e, Q'_r are equivalent linear transformations of target normalized critic networks $n_e^{Q'}, n_r^{Q'}$.

In our implementation, the scale-invariant loss functions for each of our two critic networks n_e^Q and n_r^Q are:

$$L_e^{crit} = \frac{1}{N} \sum_i \left(\frac{y_i^e - \mu_j}{\sigma_e} - n_e^Q(s, a | \theta_e^Q) \right)^2 \tag{2.7}$$

$$L_r^{crit} = \frac{1}{N} \sum_i \left(\frac{y_i^r - \mu_j}{\sigma_r} - n_r^Q(s, a, g | \theta_r^Q) \right)^2 \tag{2.8}$$

By training n_e^Q and n_r^Q to predict normalized action-value functions, we can update π_c to jointly maximize the evaluation from both the explore and exploit critics with equal importance:

$$\nabla_{\theta_c^\pi} J = \mathbb{E}_{s_t \sim \rho^\beta} [\nabla_{\theta_c^\pi} Q_c^{norm}(s_t, a, g) |_{a=\pi_c(s_t, g | \theta_c^\pi)}] \tag{2.9}$$

where

$$Q_c^{norm}(s_t, a, g) = \frac{n_e^Q(s_t, a | \theta_e^Q) + n_r^Q(s_t, a, g | \theta_r^Q)}{2}$$

Then for our three actors π_c, π_e, π_r the implemented policy gradient update rules are:

$$\nabla_{\theta_c^\pi} J \approx \frac{1}{N} \sum_i \nabla_a Q_c^{norm}(s_i, a, g) \nabla_{\theta_c^\pi} a|_{a=\pi_c(s_i, g|\theta_c^\pi)} \quad (2.10)$$

$$\nabla_{\theta_e^\pi} J \approx \frac{1}{N} \sum_i \nabla_a n_e^Q(s_i, a|\theta_e^Q) \nabla_{\theta_e^\pi} a|_{a=\pi_e(s_i|\theta_e^\pi)} \quad (2.11)$$

$$\nabla_{\theta_r^\pi} J \approx \frac{1}{N} \sum_i \nabla_a n_r^Q(s_i, a, g|\theta_r^Q) \nabla_{\theta_r^\pi} a|_{a=\pi_r(s_i, g|\theta_r^\pi)} \quad (2.12)$$

In our experiments, when we used curiosity-driven learning, we chose actions using the combined policy π_c instead of the pure explore policy π_e . The pure explore policy π_e is still useful to train the explore critic Q_e which is then used to train the combined policy π_c .

2.2 Multi-Criteria Hindsight

We define the multiple criteria in a goal as the individual target block positions that the goal specifies. In general, for other environments, criteria can be elements of a goal that require learning separate skills to accomplish. To increase the quality of data provided by hindsight experience replay, we randomly perform the hindsight goal replacement operation independently on each criteria in a goal that we are augmenting. This is done instead of replacing the entire goal with one reached later in the same episode.

Our method provides more transition samples to the agent with goals that are only partially completed later in the same episode. With normal HER, all hindsight augmented samples

Algorithm 2: Multi-Criteria HER Augmentation Step

Given:

- an augmentation probability z
- a Replay Buffer R

Sample a batch B from R

```
foreach transition sample  $(s_t||g, a_t, r_t, s_{t+1}||g)$  in  $B$  do
  foreach target block position  $p_i$  in  $g$  do
    Sample  $u \sim U(0, 1)$ 
    if  $u < z$  then
      | Sample a position  $p'_i$  that block  $i$  reached later in the same episode.
    else
      |  $p'_i \leftarrow p_i$ 
    end
  end
   $g' \leftarrow p'_0||p'_1||\dots||p'_n$ 
   $r'_t \leftarrow r(s_{t+1}, g')$ 
  replace transition sample w/  $(s_t||g', a_t, r'_t, s_{t+1}||g')$ 
end
Pass  $B$  with augmented transition samples to neural networks for training
```

Figure 2.3: Multi-Criteria HER Augmentation Step Algorithm

that the agent receives contain goals in which all criteria were satisfied at a later timestep. With multi-criteria HER, the agent will still receive a portion of goals that it later satisfied completely, and it will also receive many goals that it later only satisfied some criteria for.

In our experiments, for both binary and incremental reward formulations, using multi-criteria HER results in significant, if not critical, improvements to sample efficiency and inter-task generalization.

2.3 Curriculum

Although multi-criteria hindsight sampling allows for more sample-efficient learning and curiosity driven exploration assists in reward discovery, it was necessary to employ curriculum

learning to successfully solve multi-block stacking with sparse rewards. Training was broken into three stages, in which reaching a threshold success rate in a previous stage caused the agent to transition to the next stage. At the beginning of each stage, the DDPG algorithm was restarted, transferring only the weights of each network from a previous stage and reinitializing an empty replay buffer.

In *stage 1*, the agent trains on a non-stacking version of the block environment to help it learn fundamental skills that are transferable to the target block stacking task. The stage 1 environment is initialized with the same number of randomly placed blocks as the target stacking task. Each episode, rather than in a stack, the blocks' target positions are randomly placed on the surface of the table. A single block's target position may also be in the air instead. This stage is designed to provide less challenging tasks in which the agent can more easily discover the basic block manipulation mechanics necessary for completing the harder stacking task.

In *stage 2*, the agent trains on actual block stacking with the environment initialized at various intermediate stages of completion. At each episode, a random number of the n blocks between 0 and $n - 1$ are initialized already in the correct position on the stack. Some targets may also still be on the table rather than on the stack.

Finally, in *stage 3*, the agent trains on the target block stacking task, in which all blocks were consistently initialized on the table, away from their target locations on the stack.

Chapter 3

Experiments

In this section, we show our method’s performance on the block stacking tasks using both binary and incremental rewards. Stacking 2, 3, and 4 blocks were tested. Ablations are also shown to demonstrate the effectiveness of each our methods. We performed tests using the following configurations:

All 3: Multi-criteria HER, curiosity-driven exploration, and curriculum learning are all used with our DDPG+HER learner.

No Curiosity: Multi-criteria HER and curriculum learning are used, but all actions are chosen using π_r . An explore actor-critic and combined actor are not trained.

No Multi-Criteria: Curiosity-driven exploration and curriculum learning are used, however HER is done in the original way as defined in [1].

No Curriculum: Multi-criteria HER and curiosity-driven exploration are used, however the agent only trains on stage 3 of the curriculum, which is the actual target task of multi-block stacking.

Vanilla DDPG+HER: None of the three methods introduced in section 2 are used. This is the original DDPG+HER algorithm as in [1]. All actions are chosen using π_r .

We trained our agent using 8 to 32 parallel workers depending on the difficulty of the task. When curiosity driven-exploration was used, during experience gathering rollouts, we assigned half of the workers to take actions using π_c , and the other half using π_r . Also during experience gathering, we applied parameter-space noise [34] to the actor networks used and gaussian noise to the actions chosen. Comprehensive hyper-parameter details can be found in the supplementary materials associated with this paper.

Success rates and per-episode reward were measured during discrete testing phases in every epoch of training. During testing, actions were always chosen using π_r . An episode was considered successful if its goal g was achieved during the episode’s final state s_T .

Success rate and per-episode reward statistics were a moving average over the last 100 episodes tested on. These two statistics are shown as a function of total environment interaction timesteps for binary reward tasks in Figure 3.1 and for incremental reward tasks in Figure 3.2.

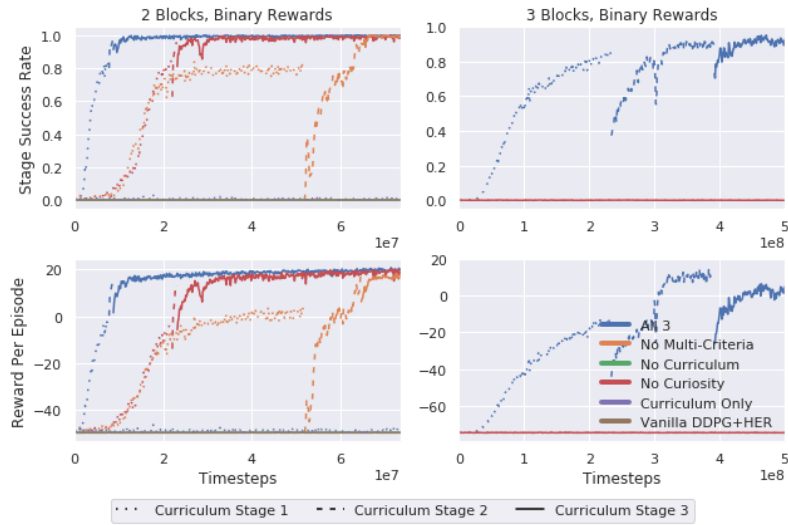


Figure 3.1: Success rates and per-episode rewards for block stacking with binary rewards. Success rates and per-episode reward values shown here are for the respective curriculum stage’s task in which they are measured.

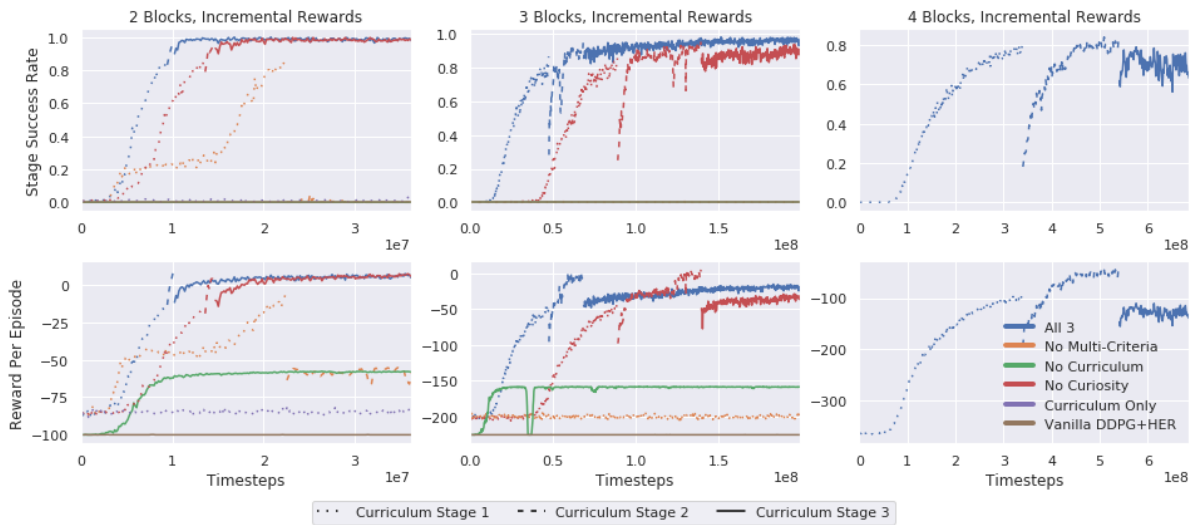


Figure 3.2: Success rates and per-episode rewards for block stacking with incremental rewards. Success rates and per-episode reward values shown here are for the respective curriculum stage’s task in which they are measured.

Table 3.1: Highest Success Rates with Binary Rewards over 100 Episode Sliding Window

Method	Stack-2	Stack-3	Stack-4
All-3	1.00	0.95	0.00
No Curiosity	1.00	0.00	-
No Multi-Criteria	1.00	0.00	-
No Curriculum	0.00	-	-
Curriculum Only	0.00	-	-
Vanilla DDPG+HER	0.00	-	-

Table 3.2: Highest Success Rates with Incremental Rewards over 100 Episode Sliding Window

Method	Stack-2	Stack-3	Stack-4
All-3	1.00	0.98	0.79
No Curiosity	0.99	0.94	-
No Multi-Criteria	0.00	0.00	-
No Curriculum	0.00	0.00	-
Curriculum Only	0.00	-	-
Vanilla DDPG+HER	0.00	-	-

Tables 3.1 and 3.2 show the highest success rates for each method on the target block stacking tasks with binary and incremental reward formulations. For methods that used a curriculum but did not reach the target task in the third stage, the final network weights were used to test performance at the target block stacking task anyways.

Vanilla DDPG+HER was unable to solve block-stacking with any number of blocks and either reward formulation.

Stacking 2 blocks with either reward formulation was solvable as long as the agent trained on the curriculum and used multi-criteria HER. Using curiosity-driven exploration without multi-criteria HER allowed the agent to make progress on stage 1 of the curriculum, but when incremental rewards were given, it failed to generalize between the stage 1 task and the stage 2 task well enough to continue learning.

Stacking 3 blocks with incremental rewards required the use of both curriculum learning and multi-criteria HER to solve. With binary rewards, stacking 3 blocks required the use of all three methods, as curiosity-driven exploration was necessary to find a reward signal.

Due to limits on computational resources, stacking 4 blocks was only tested with all three methods to measure the best possible performance. No progress was made on the binary reward environment, and in the incremental reward environment, a max success rate of 0.79 was reached on the target block stacking task.

Multi-criteria HER provided clear improvements to sample efficiency, and was necessary for stacking three or more blocks.

Agents with curiosity-driven exploration learned to solve tasks with less environment interactions than those without. With incremental rewards, block stacking was easy enough to be solved without curiosity-driven exploration, however with binary rewards, curiosity was required to solve stacking 3 blocks.

Finally, curriculum learning was necessary for any of the stacking tasks, as no method could progress on the target stacking task without first training on stages 1 and 2.

Chapter 4

Conclusion

By combining curiosity-based exploration with curriculum learning and multi-criteria HER, we are the first to solve sparse reward multi-block stacking without demonstrations. This work shows that even very challenging sparse reward environments can be solved through a combination of existing techniques. In future work, other methods of intrinsic exploration such as Go-Explore [12] might prove more effective than curiosity-driven exploration when combined with HER. In our work, we generate curricula in a hand-designed way based on domain knowledge. This might not be possible in more complex domains such as real-world robotics. Because of this, further research in automatically generating curricula is likely to be a fruitful direction when combined with HER.

Bibliography

- [1] 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 *Advances in Neural Information Processing Systems*, pages 5048–5058, 2017.
- [2] M. Bellemare, S. Srinivasan, G. Ostrovski, T. Schaul, D. Saxton, and R. Munos. Unifying count-based exploration and intrinsic motivation. In *Advances in Neural Information Processing Systems*, pages 1471–1479, 2016.
- [3] Y. Bengio, J. Louradour, R. Collobert, and J. Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48. ACM, 2009.
- [4] Y. Burda, H. Edwards, D. Pathak, A. Storkey, T. Darrell, and A. A. Efros. Large-scale study of curiosity-driven learning. *arXiv preprint arXiv:1808.04355*, 2018.
- [5] Y. Burda, H. Edwards, A. Storkey, and O. Klimov. Exploration by random network distillation. *arXiv preprint arXiv:1810.12894*, 2018.
- [6] C. Colas, O. Sigaud, and P.-Y. Oudeyer. Curious: Intrinsically motivated multi-task, multi-goal reinforcement learning. *arXiv preprint arXiv:1810.06284*, 2018.
- [7] M. P. Deisenroth and C. E. Rasmussen. Pilco: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on machine learning (ICML-11)*, pages 465–472, 2011.
- [8] M. P. Deisenroth, C. E. Rasmussen, and D. Fox. Learning to control a low-cost manipulator using data-efficient reinforcement learning. *Robotics: Science and Systems VII*, pages 78–85, 2011.
- [9] A. Deshpande and S. Sarma. Improvements on hindsight learning. *arXiv preprint arXiv:1809.06719*, 2018.
- [10] P. Dhariwal, C. Hesse, O. Klimov, A. Nichol, M. Plappert, A. Radford, J. Schulman, S. Sidor, Y. Wu, and P. Zhokhov. Openai baselines. 2018. <https://github.com/openai/baselines>.
- [11] Y. Duan, M. Andrychowicz, B. Stadie, O. J. Ho, J. Schneider, I. Sutskever, P. Abbeel, and W. Zaremba. One-shot imitation learning. In *Advances in neural information processing systems*, pages 1087–1098, 2017.

- [12] A. Ecoffet, J. Huizinga, J. Lehman, K. O. Stanley, and J. Clune. Go-explore: a new approach for hard-exploration problems. *arXiv preprint arXiv:1901.10995*, 2019.
- [13] M. Eppe, S. Magg, and S. Wermter. Curriculum goal masking for continuous deep reinforcement learning. *arXiv preprint arXiv:1809.06146*, 2018.
- [14] C. Florensa, D. Held, M. Wulfmeier, M. Zhang, and P. Abbeel. Reverse curriculum generation for reinforcement learning. *arXiv preprint arXiv:1707.05300*, 2017.
- [15] S. Forestier, Y. Mollard, and P.-Y. Oudeyer. Intrinsically motivated goal exploration processes with automatic curriculum learning. *arXiv preprint arXiv:1708.02190*, 2017.
- [16] P. Fournier, O. Sigaud, M. Chetouani, and P.-Y. Oudeyer. Accuracy-based curriculum learning in deep reinforcement learning. *arXiv preprint arXiv:1806.09614*, 2018.
- [17] G. Gordon and E. Ahissar. Hierarchical curiosity loops and active sensing. *Neural Networks*, 32:119–129, 2012.
- [18] K. Gregor, D. J. Rezende, and D. Wierstra. Variational intrinsic control. *arXiv preprint arXiv:1611.07507*, 2016.
- [19] N. Haber, D. Mrowca, L. Fei-Fei, and D. L. Yamins. Learning to play with intrinsically-motivated self-aware agents. *arXiv preprint arXiv:1802.07442*, 2018.
- [20] R. Houthoofd, X. Chen, Y. Duan, J. Schulman, F. De Turck, and P. Abbeel. Vime: Variational information maximizing exploration. In *Advances in Neural Information Processing Systems*, pages 1109–1117, 2016.
- [21] R. Kaushik, K. Chatzilygeroudis, and J.-B. Mouret. Multi-objective model-based policy search for data-efficient learning with sparse rewards. *arXiv preprint arXiv:1806.09351*, 2018.
- [22] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [23] A. S. Klyubin, D. Polani, and C. L. Nehaniv. Empowerment: A universal agent-centric measure of control. In *Evolutionary Computation, 2005. The 2005 IEEE Congress on*, volume 1, pages 128–135. IEEE, 2005.
- [24] A. Laversanne-Finot, A. Péré, and P.-Y. Oudeyer. Curiosity driven exploration of learned disentangled goal spaces. *arXiv preprint arXiv:1807.01521*, 2018.
- [25] 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. *arXiv preprint arXiv:1509.02971*, 2015.
- [26] D. Y.-J. Little and F. T. Sommer. Learning and exploration in action-perception loops. *Frontiers in neural circuits*, 7:37, 2013.

- [27] S. McAleer, F. Agostinelli, A. Shmakov, and P. Baldi. Solving the rubik’s cube with approximate policy iteration. In *International Conference on Learning Representations*, 2019.
- [28] S. Mohamed and D. J. Rezende. Variational information maximisation for intrinsically motivated reinforcement learning. In *Advances in neural information processing systems*, pages 2125–2133, 2015.
- [29] A. Nair, B. McGrew, M. Andrychowicz, W. Zaremba, and P. Abbeel. Overcoming exploration in reinforcement learning with demonstrations. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6292–6299. IEEE, 2018.
- [30] G. Ostrovski, M. G. Bellemare, A. v. d. Oord, and R. Munos. Count-based exploration with neural density models. *arXiv preprint arXiv:1703.01310*, 2017.
- [31] D. Pathak, P. Agrawal, A. A. Efros, and T. Darrell. Curiosity-driven exploration by self-supervised prediction. In *International Conference on Machine Learning (ICML)*, volume 2017, 2017.
- [32] A. Péré, S. Forestier, O. Sigaud, and P.-Y. Oudeyer. Unsupervised learning of goal spaces for intrinsically motivated goal exploration. *arXiv preprint arXiv:1803.00781*, 2018.
- [33] M. Plappert, M. Andrychowicz, A. Ray, B. McGrew, B. Baker, G. Powell, J. Schneider, J. Tobin, M. Chociej, P. Welinder, et al. Multi-goal reinforcement learning: Challenging robotics environments and request for research. *arXiv preprint arXiv:1802.09464*, 2018.
- [34] M. Plappert, R. Houthoofd, P. Dhariwal, S. Sidor, R. Y. Chen, X. Chen, T. Asfour, P. Abbeel, and M. Andrychowicz. Parameter space noise for exploration. *arXiv preprint arXiv:1706.01905*, 2017.
- [35] I. Popov, N. Heess, T. Lillicrap, R. Hafner, G. Barth-Maron, M. Vecerik, T. Lampe, Y. Tassa, T. Erez, and M. Riedmiller. Data-efficient deep reinforcement learning for dexterous manipulation. *arXiv preprint arXiv:1704.03073*, 2017.
- [36] P. Rauber, A. Ummadisingu, F. Mutz, and J. Schmidhuber. Hindsight policy gradients. *arXiv preprint arXiv:1711.06006*, 2017.
- [37] M. Riedmiller, R. Hafner, T. Lampe, M. Neunert, J. Degraeve, T. Van de Wiele, V. Mnih, N. Heess, and J. T. Springenberg. Learning by playing-solving sparse reward tasks from scratch. *arXiv preprint arXiv:1802.10567*, 2018.
- [38] J. Schmidhuber. A possibility for implementing curiosity and boredom in model-building neural controllers. In *Proc. of the international conference on simulation of adaptive behavior: From animals to animats*, pages 222–227, 1991.
- [39] J. Schmidhuber. Formal theory of creativity, fun, and intrinsic motivation (1990–2010). *IEEE Transactions on Autonomous Mental Development*, 2(3):230–247, 2010.

- [40] H. P. van Hasselt, A. Guez, M. Hessel, V. Mnih, and D. Silver. Learning values across many orders of magnitude. In *Advances in Neural Information Processing Systems*, pages 4287–4295, 2016.
- [41] M. Večerík, T. Hester, J. Scholz, F. Wang, O. Pietquin, B. Piot, N. Heess, T. Rothörl, T. Lampe, and M. Riedmiller. Leveraging demonstrations for deep reinforcement learning on robotics problems with sparse rewards. *arXiv preprint arXiv:1707.08817*, 2017.
- [42] R. Zhao and V. Tresp. Curiosity-driven experience prioritization via density estimation. 2018.
- [43] R. Zhao and V. Tresp. Energy-based hindsight experience prioritization. *arXiv preprint arXiv:1810.01363*, 2018.

Appendices

A Links

A video showcasing this project is available at

<https://youtu.be/stZX4o0H8Ro>

Code for our modified DDPG Learner is available at

<https://github.com/CDMCH/ddpg-with-curiosity-and-multi-criteria-her>

and code for our block stacking environments is available at

<https://github.com/CDMCH/gym-fetch-stack>

Our DDPG learner uses code modified from the OpenAI baselines repository [10].

B Experiment Details

Observation and goal network inputs were normalized to have a mean of zero and standard deviation of one. Once normalized, they were also clipped to the range $[-5, 5]$.

All networks were fully connected with 3 hidden layers and 256 hidden units in each layer. Hidden layers used ReLU activations, while the output layers of actor networks used tanh. The action space was re-scaled to fit the tanh range of $[-1, 1]$, and to prevent vanishing gradients, the preactivations of the actor output layers were penalized by the square of their magnitude with a coefficient of 0.001.

The DDPG algorithm was run in parallel using multiple message passing interface (MPI)-based workers. Network parameters and normalization statistics were averaged across workers during update steps. The actor policy, π_e , π_r , or π_c that each worker used during experience gathering was set as a hyperparameter. All workers used π_r during performance testing. Different worker amounts were used depending on the difficulty of the task:

Table B.1: Parallel Worker Amounts by Task

Task	Number of MPI Workers
Stack 2, Sparse Rewards	8
Stack 3, Sparse Rewards	32
Stack 4, Sparse Rewards	32
Stack 2, Incremental Rewards	8
Stack 3, Incremental Rewards	8
Stack 4, Incremental Rewards	32

The following hyperparameters were used in our experiments:

Table B.2: Hyperparameters for Block Stacking Tasks

Hyperparameter	Value
Optimizer	Adam [22]
n_r Learning Rate	0.001
n_r L2 Regularization Coefficient	0
π_r Learning Rate	0.001
Target Exploit Actor-Critic Polyak-averaging Coefficient	0.001
n_e Learning Rate	0.001
n_e L2 Regularization Coefficient	0.01
π_e Learning Rate	0.001
Target Explore Actor-Critic Polyak-averaging Coefficient:	0.05
π_c Learning Rate	0.001
π_c Explore vs Exploit Critic Weighting	0.5, 0.5
D Learning Rate	0.007
Episode Time Horizon	50 * num blocks
γ	1 - 1/episode time horizon
MPI Worker Replay Buffer Size	10^6 transitions
Parameter Space Noise σ Target	0.1
Gaussian Action Noise σ	0.04
Traditional HER Augmentation Probability (when used)	0.8
Multi-Criteria HER Augmentation Probability (when used)	0.8
Cycles per Epoch	50
Experience Gathering Episodes per Cycle	8 (per MPI worker)
Training Batches per cycle	8
Network Update Batch Size	1024 transitions (per MPI worker)
Test Episode Rollouts Per Epoch	50