
Guided Exploration with Proximal Policy Optimization using a Single Demonstration

Gabriele Libardi¹ Sebastian Dittert¹ Gianni De Fabritiis^{1,2}

Abstract

Solving sparse reward tasks through exploration is one of the major challenges in deep reinforcement learning, especially in three-dimensional, partially-observable environments. Critically, the algorithm proposed in this article is capable of using a single human demonstration to solve hard-exploration problems. We train an agent on a combination of demonstrations and own experience to solve problems with variable initial conditions and we integrate it with proximal policy optimization (PPO). The agent is also able to increase its performance and to tackle harder problems by replaying its own past trajectories prioritizing them based on the obtained reward and the maximum value of the trajectory. We finally compare variations of this algorithm to different imitation learning algorithms on a set of hard-exploration tasks in the Animal-AI Olympics environment. To the best of our knowledge, learning a task in a three-dimensional environment with comparable difficulty has never been considered before using only one human demonstration.

1. Introduction

Exploration is one of the most challenging problems in reinforcement learning. Although significant progress has been made in solving this problem in recent years (Badia et al., 2020a;b; Burda et al., 2018; Pathak et al., 2017; Ecoffet et al., 2019), many of the solutions rely on very strong assumptions such as access to the agent position and deterministic, fully observable or low dimensional environments. For three-dimensional, partially observable stochastic environments with no access to the agent position the exploration problem still remains unsolved.

¹Computational Science Laboratory, Universitat Pompeu Fabra (UPF) ²ICREA . Correspondence to: Gabriele Libardi <gabrielelibardi@yahoo.it>, Gianni De Fabritiis <gianni.defabritiis@upf.edu>.

Learning from demonstrations allows to directly bypass this problem but it only works under specific conditions, e.g. large number of demonstration trajectories or access to an expert to query for optimal actions (Ross et al., 2011). Furthermore, a policy learned from demonstrations in this way will only be optimal in the vicinity of the demonstrator trajectories and only for the initial conditions of such trajectories. Our work is at the intersection of these two classes of solutions, exploration and imitation, in that we use only one trajectory from the demonstrator per problem to solve hard-exploration tasks^{1,2}.

This approach has been investigated before by Paine et al. (2019) (for a thorough comparison see Section 5.2). We propose the first implementation based on on-policy algorithms, in particular PPO. Henceforth we refer to the version of the algorithm we put forward as PPO + Demonstrations (PPO+D). Our contributions can be summarized as follows:

1. In our approach, we treat the demonstrations trajectories as if they were the result of lucky actions taken by the agent in the real environment. We can do this because in PPO the policy only specifies a distribution over the action space. We force the actions of the policy to equal the demonstration actions instead of sampling from the policy distribution and in this way we accelerate the exploration phase. We use importance sampling to account for sampling from a distribution different than the policy. The frequency with which the demonstrations are sampled depends on an adjustable hyperparameter ρ , as described in Paine et al. (2019).
2. Our algorithm includes the successful trajectories in the replay buffer during training and treats them as demonstrations.
3. The non-successful trajectories are ranked according to their maximum estimated value and are stored in a different replay buffer.

¹A video showing the experimental results is available at <https://www.youtube.com/playlist?list=PLBeSdcnDP2WFQWLBrLGSkwtitneOelcm->

²The code is available here: https://github.com/compsciencelab/ppo_D

4. We mitigate the effect of catastrophic forgetting by using the maximum reward and the maximum estimated value of the trajectories to prioritize experience replay.

PPO+D is only in part on-policy as a fraction of its experience comes from a replay buffer and therefore was collected by an older version of the policy. The importance sampling is limited to the action loss in PPO and does not adjust the target value in the value loss as in [Espeholt et al. \(2018\)](#).

We found that this new algorithm is capable of solving problems that are not solvable using normal PPO, behavioral cloning, GAIL, nor combining behavioral cloning and PPO. PPO+D is very easy to implement by only slightly modifying the PPO algorithm. Crucially, the learned policy is significantly different and more efficient than the demonstration used in the training.

To test this new algorithm we created a benchmark of hard-exploration problems of varying levels of difficulty using the Animal-AI Olympics challenge environment ([Beyret et al., 2019](#); [Crosby et al., 2019](#)). All the tasks considered have random initial position and the PPO policy uses entropy regularization so that memorizing the sequence of actions of the demonstration will not suffice to complete any of the tasks. We also test PPO+D on a modified sparse version of the ReacherPyBulletEnv-v0 and LunarLander-v2 tasks based on the implementations available in [Brockman et al. \(2016\)](#); [Coumans & Bai \(2017\)](#).

2. Related work

Different attempts have been made to use demonstrations efficiently in hard-exploration problems. In [Salimans & Chen \(2018\)](#) the agent is able to learn a robust policy only using one demonstration. The demonstration is replayed for n steps after which the agent is left to learn on its own. By incrementally decreasing the number of steps n , the agent learns a policy that is robust to randomness (introduced in this case by using sticky actions or no-ops ([Machado et al., 2018](#)), since the game is fundamentally a deterministic one). However, this approach only works in a fully deterministic environment since replaying the actions has the role of resetting the environment to a particular configuration. This method of resetting is obviously not feasible in a stochastic environment.

[Ecoffet et al. \(2019\)](#) introduces another algorithm that largely exploits the determinism of the environment by resetting to previously reached states. It works by maximizing the diversity of the states reached. It is able to identify such diversity among states by down-sampling the observations, and by considering as different states only those observations that have a different down-sampled image. This works remarkably well in two dimensional environments, but is

unlikely to work in three-dimensional, stochastic environments.

Self-supervised prediction approaches, such as [Pathak et al. \(2017\)](#); [Burda et al. \(2018\)](#); [Schmidhuber \(2010\)](#); [Badia et al. \(2020b\)](#), have been used successfully in stochastic environments, although they are less effective in three-dimensional environments. Another class of algorithms designed to solve exploration problems are count-based methods ([Tang et al., 2017](#)). These algorithms keep track of the states where the agent has been before (if we are dealing with a prohibitive number of states, the dimensions along which the agent moves can be discretized), and give the agent an incentive (in the form of a bonus reward) for visiting new states. This approach assumes we have a reliable way to track the position of the agent. This assumption is removed in the case of pseudo-counts when using neural density models ([Bellemare et al., 2016](#); [Ostrovski et al., 2017](#)). [Choi et al. \(2018\)](#) proposes an alternative method for localizing the agent based on an attention mask over pixels in the image used in predicting the action the agent took. Like [Pathak et al. \(2017\)](#) we think this solution too is probably biased towards two dimensional environments.

An empirical comparison between these two classes of exploration algorithms was made in ([Baker et al., 2019](#)), where agents compete against each other leveraging the use of tools that they learn to manipulate. They found the count-based approach works better when applied not only to the agent position, but also to relevant entities in the environment (such as the positions of objects). When only given access to the agent position, the Random Network Distillation algorithm ([Burda et al., 2018](#)) was found to lead to a higher performance.

Some other works focus on leveraging the use of expert demonstrations while still maximizing the reward. These methods allow the agent to outperform the expert demonstration in many problems. [Hester et al. \(2018\)](#) combines temporal difference updates in the Deep Q-Network (DQN) algorithm with supervised classification of the demonstrator’s actions. [Kang et al. \(2018\)](#) proposes to effectively leverage available demonstrations to guide exploration through enforcing occupancy measure matching between the learned policy and current demonstrations.

Other approaches, such as [Duan et al. \(2017\)](#); [Zhou et al. \(2019\)](#) pursue a meta-learning strategy where the agent learns to learn from demonstrations, such approaches are perhaps the most promising, but they require at least a demonstration for each task for a subset of all tasks.

Generative adversarial imitation learning (GAIL) ([Ho & Ermon, 2016](#)) has never been successfully applied to complex partially observable environments that require memory ([Paine et al., 2019](#)). InfoGAIL ([Li et al., 2017](#)) has been used

to learn from images but, unlike in this work the policy does not require recurrence to complete the tasks. Behavioral Cloning (BC) is the most basic imitation learning technique, it is equivalent to supervised classification of the demonstrator’s actions (Rahmatizadeh et al., 2018). Both GAIL and BC have been used in the Obstacle Tower Challenge (Juliani et al., 2019), but are alone insufficient for solving hard-exploration problems (Alex).

Oh et al. (2018) present an off-policy algorithm that learns to reproduce the agent’s past good decisions. Their work focuses mainly on advantage actor-critic, but the idea was tested also with PPO. Instead, PPO+D utilizes importance sampling, leverages expert (human) demonstrations, uses prioritized experience replay based on the maximum value of each trajectory to overcome catastrophic forgetting and the ratio of demonstrations replayed during training can be explicitly controlled.

3. Demonstration guided exploration with PPO+D

Our approach attempts to exploit the idea of combining demonstrations with the agent’s own experience in an on-policy algorithm such as PPO. This approach is particularly effective for hard-exploration problems. One can view the replay of demonstrations as a possible trajectory of the agent in the current environment. This means that the only point we are interfering with the classical PPO is when sampling, which is substituted by simply replaying the demonstrations. A crucial difference between PPO+D and R2D3 (Paine et al., 2019) is that we do consider sequences that contain entire episodes in them, and therefore using recurrence is much more straightforward. From the perspective of the agent it is as if it is always lucky when sampling the actions, and in doing so it is skipping the exploration phase. The importance sampling formula provides the average rewards over policy $\pi_{\theta'}$ given trajectories generated from a different policy $\pi_{\theta}(a|s)$:

$$\mathbb{E}_{\pi_{\theta'}}[r_t] = \mathbb{E}_{\pi_{\theta}} \left[\frac{\pi_{\theta'}(a_t|s_t)}{\pi_{\theta}(a_t|s_t)} r_t \right], \quad (1)$$

where $r_t \sim R(r|a_t, s_t)$ is the environment reward given state s_t and action a_t at time t . $\mathbb{E}_{\pi_{\theta}}$ indicates that the average is over trajectories drawn from the parameterized policy π_{θ} . The importance sampling term $\frac{\pi_{\theta'}(a_t|s_t)}{\pi_{\theta}(a_t|s_t)}$ accounts for the correction in the change of the distribution over actions in the policy $\pi_{\theta'}(a_t|s_t)$. By maximizing $\mathbb{E}_{\pi_{\theta}} \left[\frac{\pi_{\theta'}(a_t|s_t)}{\pi_{\theta}(a_t|s_t)} r_t \right]$ over the parameters θ' a new policy is obtained that is on average better than the old one. The PPO loss function is then defined as

$$L_t^{PPO}(\theta) = \mathbb{E}_t \left[L_t^{CLIP}(\theta) + c_1 L_t^{VF}(\theta) + c_2 S^{\pi_{\theta}}(s_t, a_t) \right], \quad (2)$$

where

$$L_t^{CLIP}(s_t, a_t, \theta', \theta) = \min \left(\frac{\pi_{\theta'}(a_t|s_t)}{\pi_{\theta}(a_t|s_t)} A^{\pi_{\theta}}(s_t, a_t), g(\epsilon, A^{\pi_{\theta}}(s_t, a_t)) \right) \quad (3)$$

and c_1 and c_2 are coefficients, L_t^{VF} is the squared-error loss $(V_{\theta}(s_t) - V_t^{arg})^2$, $A^{\pi_{\theta}}$ is an estimate of the advantage function and S is the entropy of the policy distribution. The entropy term is designed to help keep the search alive by preventing convergence to a single choice of output, especially when several choices all lead to roughly the same reward (Williams & Peng, 1991).

Let \mathcal{D} be the set of trajectories $\tau = (s_0, a_0, s_1, a_1, \dots)$ for which we have a demonstration, then $\pi_{\mathcal{D}}(a_t|s_t) = 1$ for (a_t, s_t) in \mathcal{D} and 0 otherwise. This is a policy that only maps observations coming from the demonstration buffer to a distribution over actions. Such distribution assigns probability one to the action taken in the demonstration and zero to all the remaining actions. The algorithm decides where to sample trajectories from each time an episode is completed (for running out of time or for completing the task). We define \mathcal{D} to be the set of all trajectories that can get replayed at any given time $\mathcal{D} = \mathcal{D}_V \cup \mathcal{D}_R$, where \mathcal{D}_V are the trajectories collected prioritizing the value estimate (V stands for value), and \mathcal{D}_R (R stands for reward) contains the initial human demonstration and successful trajectories the agent collects. Since the sparse environments we tested PPO+D on only have one source of reward, any trajectory can either be considered successful in completing the task or unsuccessful. The agent samples from the trajectories in \mathcal{D}_R with probability ρ , \mathcal{D}_V with probability ϕ and from the real environment Env with probability $1 - \rho - \phi$ subject to $\rho + \phi \leq 1$.

The behavior of the policy can be defined as follows:

$$\pi_{\theta}^{\phi, \rho} = \begin{cases} \pi_{\mathcal{D}_R}, & \text{if sampled from } \mathcal{D}_R \\ \pi_{\mathcal{D}_V}, & \text{if sampled from } \mathcal{D}_V \\ \pi_{\theta}, & \text{if sampled from Env} \end{cases}$$

In PPO+D we substitute the current policy π_{θ} with the policy $\pi_{\theta}^{\phi, \rho}$, since this is the policy used to collect the demonstration trajectories. The clipping term is then changed to

correct for it,

$$L_t^{PPO+D}(s_t, a_t, \theta', \theta) = \min \left(\frac{\pi_{\theta'}(a_t|s_t)}{\pi_{\theta}^{\phi, \rho}(a_t|s_t)} A^{\pi_{\theta}^{\phi, \rho}}(s_t, a_t), g(\epsilon, A^{\pi_{\theta}^{\phi, \rho}}(s_t, a_t)) \right). \quad (4)$$

3.1. Self-imitation and prioritized experience replay

We hold the total size of the buffers $|\mathcal{D}|$ constant throughout the training. At the beginning of the training we only provide the agent with one demonstration trajectory, so $|\mathcal{D}_R| = 1$ and $|\mathcal{D}_V| = |\mathcal{D}| - 1$ as $|\mathcal{D}| = |\mathcal{D}_V| + |\mathcal{D}_R|$. \mathcal{D}_V is only needed at the beginning to collect the first successful trajectories. As $|\mathcal{D}_R|$ increases and we have more variety of trajectories in \mathcal{D}_R , we decrease the probability ϕ of replaying trajectories from \mathcal{D}_V . After $|\mathcal{D}_R|$ is large enough, replaying the trajectories from \mathcal{D}_V is more of a hindrance. The main reason for treating these two buffers separately is to slowly anneal from one to the other avoiding the hindrance in the second phase.

$|\mathcal{D}_V|$ and $|\mathcal{D}_R|$ are the quantities that are annealed according to the following formulas, each time a new successful trajectory is found and is added to \mathcal{D}_R :

$$\begin{aligned} \rho &= \rho + \frac{\phi_0}{|\mathcal{D}_V|_0}; & \phi &= \phi - \frac{\phi_0}{|\mathcal{D}_V|_0}; \\ |\mathcal{D}_V| &= |\mathcal{D}_V| - 1; & |\mathcal{D}_R| &= |\mathcal{D}_R| + 1 \end{aligned} \quad (5)$$

where $|\mathcal{D}_R|_0$ and $|\mathcal{D}_V|_0$ are the maximum size respectively for \mathcal{D}_R and \mathcal{D}_V and ϕ_0 is the value of ϕ at the beginning of the training. We define the probability of sampling trajectory τ_i as $P(i) = \frac{p_i^\alpha}{\sum_k p_k^\alpha}$, where $p_i = \max_t V_\theta(s_t)$, and α is a hyperparameter. We shift the value of p_i for all the trajectories in \mathcal{D}_V so as to guarantee $p_i \geq 0$. We only keep up to $|\mathcal{D}_V|$ unsuccessful trajectories at any given time.

Values are only updated for the replayed transitions. Successful trajectories are similarly sampled from \mathcal{D}_R with a uniform distribution (a better strategy could be to sample according to their length), and the buffer is updated following a FIFO strategy. We introduced the value-based experience replay because we get to a level of complexity in some tasks that we could not solve by using self-imitation solely based on the reward. These are the tasks that the agent has trouble solving even once because the sequence of actions is too long or complicated. We prefer the value estimate as a selection criteria rather than the more commonly used TD-error as we speculate it is more effective at retaining promising trajectories over long periods of time. Crucially for the unsuccessful trajectories it is possible for the maximum value estimate not to be zero when the observations are similar to the ones seen in the demonstration.

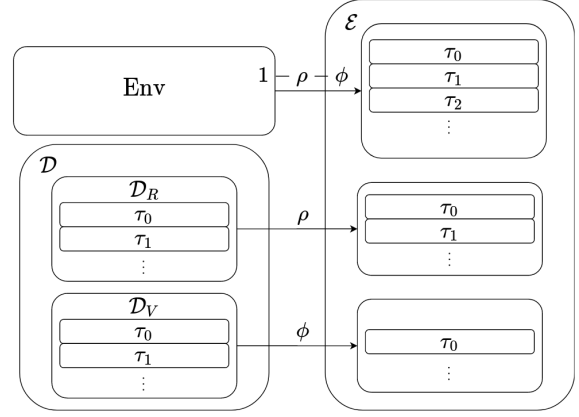


Figure 1. **Learner diagram.** The PPO+D learner samples batches that are a mixture of demonstrations and the experiences the agent collects by interacting with the environment.

Algorithm 1 PPO+D

```

Initialize parameters  $\theta$ 
Initialize replay buffer  $\mathcal{D}_V \leftarrow \{\}$ 
Initialize replay buffer  $\mathcal{D}_R \leftarrow \{\tau_{\mathcal{D}}\}$ 
Initialize rollout storage  $\mathcal{E} \leftarrow \{\}$ 
for every update do
  for actors  $1, 2, \dots, N$  do
    Sample  $\tau$  from  $\{\mathcal{D}_V, \mathcal{D}_R, \text{Env}\}$  based on  $\rho$  and  $\phi$ 
    if  $\tau \in \mathcal{D}_R$  then
      for steps  $1, 2, \dots, T$  do
        Replay a transition from
        buffer  $s_t, a_t, r_t, s_{t+1} \sim \pi_{\mathcal{D}_R}(a_t|s_t)$ 
        Store transition  $\mathcal{E} \leftarrow \mathcal{E} \cup \{(s_t, a_t, r_t)\}$ 
      end for
    else if  $\tau \in \mathcal{D}_V$  then
      for steps  $1, 2, \dots, T$  do
        Replay a transition from buffer
         $s_t, a_t, r_t, s_{t+1} \sim \pi_{\mathcal{D}_V}(a_t|s_t)$ 
        Store transition  $\mathcal{E} \leftarrow \mathcal{E} \cup \{(s_t, a_t, r_t)\}$ 
      end for
    else if  $\tau \in \text{Env}$  then
      for steps  $1, 2, \dots, T$  do
        Execute an action in the environment
         $s_t, a_t, r_t, s_{t+1} \sim \pi_\theta(a_t|s_t)$ 
        Store transition  $\mathcal{E} \leftarrow \mathcal{E} \cup \{(s_t, a_t, r_t)\}$ 
      end for
    end if
    Compute advantages estimates  $\hat{A}_1, \hat{A}_2, \dots, \hat{A}_T$ 
  end for
  Optimize  $L^{PPO}$  wrt  $\theta$ , with  $K$  epochs and
  minibatches size  $M \leq NT$ 
   $\theta \leftarrow \theta - \eta \nabla_\theta L^{PPO}$ 
  Empty rollout storage  $\mathcal{E} \leftarrow \{\}$ 
  Update  $|\mathcal{D}_V|, |\mathcal{D}_R|, \rho$  and  $\phi$  according to Equations 5
end for

```

We think that this plays a role in countering the effects of catastrophic forgetting, thus allowing the agent to combine separately learned sub-behaviors in a successful policy. We illustrate this with some example trajectories of the agent shown in the Appendix in Figure 8. The value estimate is noisy and as a consequence of that, trajectories that have a low maximum value estimate on first visits may not be replayed for a long time or never as pointed out in (Schaul et al., 2015). However, for our strategy to work it is enough for some of the promising trajectories to be collected and replayed by this mechanism.

4. Experiments

4.1. The Animal-AI Olympics challenge environment

The recent successes of deep reinforcement learning (DRL) (Mnih et al., 2015; Silver et al., 2017; Schulman et al., 2017; Schrittwieser et al., 2019; Srinivas et al., 2020) have shown the potential of this field, but at the same time have revealed the inadequacies of using games (such as the ATARI games (Bellemare et al., 2013)) as a benchmark for intelligence. These inadequacies have motivated the design of more complex environments that will provide a better measure of intelligence.

The Obstacle Tower Environment (Juliani et al., 2019), the Animal AI Olympics (Crosby et al., 2019), the Hard-Eight Task Suite (Paine et al., 2019) and the DeepMind Lab (Beattie et al., 2016) all exhibit sparse rewards, partial observability and highly variable initial conditions. For some tests we use the Animal-AI Olympics challenge environment. The aim of the Animal-AI Olympics is to translate animal cognition into a benchmark of cognitive AI (Crosby et al., 2019).

The environment contains an agent enclosed in a fixed size arena. Objects can spawn in this arena, including positive and negative rewards (green, yellow and red spheres) that the agent must obtain or avoid. This environment has basic physics rules and a set of objects such as food, walls, negative-reward zones, movable blocks and more. The playground can be configured by the participants and they can spawn any combination of objects in preset or random positions. We take advantage of the great flexibility allowed by this environment to design hard-exploration problems for our experiments.

We also included some experiments on tasks that have already been extensively studied in the literature: the ReacherPyBulletEnv-v0 (Coumans & Bai, 2017) and LunarLander-v2 (Brockman et al., 2016) tasks. However, because the use of PPO+D is only justified when dealing with very sparse reward environments we modified the reward function of these tasks to fulfill this criteria.

4.2. Experimental setting

We designed our experiments in order to answer the following questions: Can the agent learn a policy in a non-deterministic hard-exploration environment only with one human demonstration? Is the agent able to use the human demonstration to learn to solve a problem with different initial conditions (agent and boxes initial positions) than the demonstration trajectory? How does the hyperparameter ϕ influence the performance during training? The four tasks in the Animal AI Olympics environments we used to evaluate the agent are described as follows:

- **One box easy**

The agent has to move a single box, always spawned at the same position, in order to bridge a gap and be able to access the reward once it climbs up the ramp (visible in pink). The agent can be spawned in the range ($X : 0.5 - 39.5, Y : 0.5 - 39.5$) if an object is not already present at the same location (Fig. 2a).

- **One box hard**

The agent has to move a single box in order to bridge a gap and be able to access the reward, this time two boxes are spawned at any of four positions $A : (X : 10, Y : 10), B : (X : 10, Y : 30), C : (X : 30, Y : 10), D : (X : 30, Y : 30)$. The agent can be spawned in the range ($X : 0.5 - 39.5, Y : 0.5 - 39.5$) if an object is not already present at the same location (Fig. 2b).

- **Two boxes easy**

The agent has to move two boxes in order to bridge a larger gap and be able to access the reward, this time two boxes are spawned at any of four positions $A : (X : 10, Y : 10), B : (X : 10, Y : 30), C : (X : 30, Y : 10), D : (X : 30, Y : 30)$. The agent can be spawned in the range ($X : 15.0 - 20.0, Y : 0.5 - 15.0$) if an object is not already present at the same location (Fig. 2c).

- **Two boxes hard**

The agent has to move two boxes in order to bridge a larger gap and be able to access the reward. Two boxes are spawned at two fixed positions $A : (X : 10, Y : 10), B : (X : 10, Y : 30)$. A wall acts as a barrier in the middle of the gap, to prevent the agent from "surfing" a single box. The agent can be spawned in the range ($X : 15.0 - 20.0, Y : 5.0 - 10.0$) if an object is not already present at the same location (Fig. 2d).

For the Sparse LunarLander-v2 tasks the reward has zero value if the agent is unable to land and one otherwise. For the ReacherPyBulletEnv-v0 task the common reward function contains a term proportional to the negative euclidean distance between the tip of the robot arm and the target position. In the Sparse ReacherPyBulletEnv-v0 the reward is modified as follows:

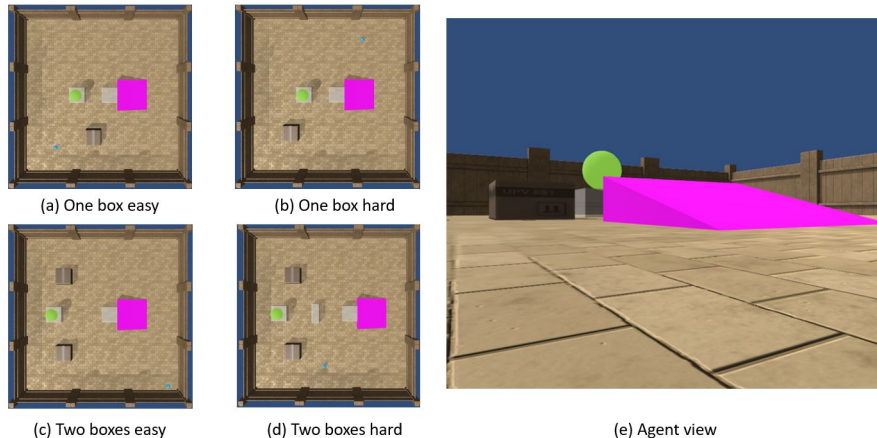


Figure 2. **Tasks.** In each of the tasks there is only one source of reward and the position of some of the objects is random, so each episode is different. The agent has no access to the aerial view, instead it partially observes the world through a first person view of the environment. All of the tasks are either inspired or adapted from the test set in the Animal-AI Olympics competition. A video showing the experimental results is available at <https://www.youtube.com/playlist?list=PLBeSdcnDP2WFQWLBRLGskwtitneOelcm-> .

$$r = \begin{cases} T - d, & \text{if } d \leq T \\ 0, & \text{if } d > T \end{cases}$$

with $T \geq 0$ and $d \geq 0$, where d is the euclidean distance between tip of the arm and target position, and T is a threshold value. For our experiments we chose $T = 1$.

5. Results

5.1. Comparison with baselines and generalization

In Figure 4 we compare PPO+D with parameters $\rho = 0.1, \phi = 0.3$ to the behavioral cloning baselines (with 100 and 10 demonstrations), to GAIL (with 100 demonstrations) and with PPO+BC. PPO+BC combines PPO and BC in an a similar way to PPO+D: with probability ρ a sample is drawn from the demonstrations and the policy is updated using the BC loss (see Section A in the Appendix for more details). The value loss function remains unchanged during the BC update.

We tested the GAIL implementation on a simple problem to validate it (see Section C in the Appendix). For behavioral cloning we trained for 3000 learner steps (updates of the policy) with learning rate 10^{-5} . It is clear that PPO+D is able to outperform the baseline in all four problems. The performance of PPO+D varies considerably from task to task. This reflects the considerable difference in the range of the initial conditions for different tasks. In the tasks "One box easy" and "Two boxes hard" only the position of the agent sets different instances of the task apart. The initial positions of the boxes only play a role in the tasks "One box hard" and "Two boxes easy". Under closer inspection we could verify that in these two tasks the policy fails to

generalize to configurations of boxes that are not seen in the demonstration, but does generalize well in all tasks for very different agent starting positions and orientations.

This could be because there is a smooth transition between some of the initial conditions. Due to this fact, if the agent is able to generalize even only between initial conditions that are close, it will be able to progressively learn to perform well for all initial conditions starting from one demonstration. In other words the agent is automatically using a curriculum learning strategy, starting from the initial conditions that are closer to the demonstration. This approach fails when there is an abrupt transition between initial conditions, such as for different boxes configurations.

During training we realized that the policy exploited the 3D engine in the task "Two boxes easy" as it managed to "surf" one of the boxes, in this way avoiding to move the remaining box (a similar behavior was reported in (Baker et al., 2019)). To avoid this trick and make sure the agent is capable of such more complex behavior we introduced "Two boxes hard". We decided to reduce the range of the initial conditions in this last task, as we already verified that the agent can learn from such variable conditions in tasks "One box hard" and "Two boxes easy". This last experiment only tests the agent for a more complex behavior. In the tasks "One box hard" and "Two boxes easy" the agent could achieve higher performance given more training.

In Figure 3 we compare PPO+D with PPO+BC and vanilla PPO on the tasks Sparse LunarLander-v2 and Sparse ReacherPyBulletEnv-v0. These tasks are considerably easier than the Animal AI Olympics tasks but are useful to underlie the difference in performance between PPO+D and PPO+BC. For these experiments we used ten demonstra-

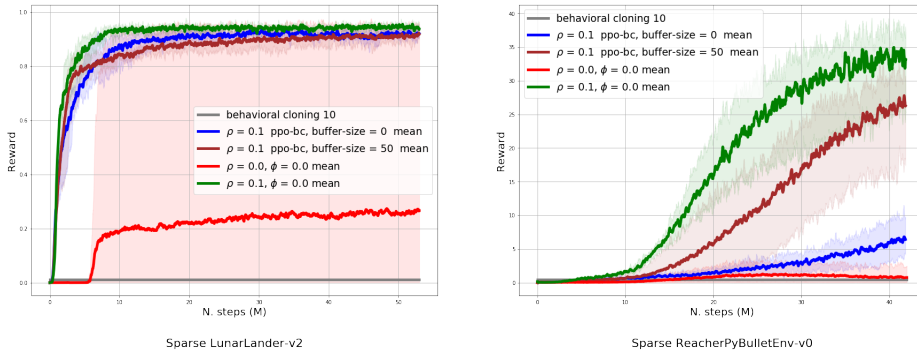


Figure 3. **Experiments.** Performance of PPO+D (green), PPO+BC with reward buffer (buffer size =50, brown) and without it (buffer size =0, blue) vanilla PPO (red) and normal behavior cloning (gray). For Sparse ReacherPyBulletEnv-v0 we chose the hyperparameters $\rho = 0.1, \phi = 0.0$ and for Sparse LunarLander-v2 $\rho = 0.3, \phi = 0.0$, which we found to be optimal for both PPO+D and PPO+BC. For both experiments the agent has ten human demonstrations available to learn from. The curves show the mean min and max reward for each of the baselines across 5 different seeds.




Figure 4. **Experiments.** Performance of behavioral cloning with ten and a hundred recorded human demonstrations and PPO+D with $\rho = 0.1, \phi = 0.3$ and just one demonstration. The curves represent the mean, min and max performance for each of the baselines across 5 different seeds, except for vanilla PPO, GAIL and PPO+BC (3 seeds) where the reward is consistently close to zero. The BC agent sporadically obtains some rewards. GAIL with a hundred demonstrations never achieves any reward. PPO+BC has only access to one demonstration, like PPO+D. It occasionally solves the task but it is unable to achieve high performance.

tions. For both tasks we first tuned the ρ hyperparameter for both PPO+D and PPO+BC. For PPO+BC we consider two versions, one in which it keeps adding successful trajectory to a replay buffer in the same way PPO+D does, in the other it does not have such replay buffer. In these two tasks we do not employ the value buffer as this is of any use only in very hard exploration tasks. Surprisingly even vanilla PPO is able to solve Sparse LunarLander-v2 sporadically (in one seed over five), although its performance exhibits high variance. Both variants of PPO+BC successfully learn the task, the difference with PPO+D is small but consistent over all seeds. In the more challenging task Sparse ReacherPyBulletEnv-v0 the differences in performance between PPO+D and PPO+BC with and without reward buffer are considerably amplified.

We emphasise that PPO+D is designed to perform well on hard-exploration problems with stochastic environment and variable different conditions, we tested it on the of the Atari

environment BreakoutNoFrameskip-v4 and conclude that it does not lead to better performance than vanilla PPO when the reward is dense. PPO+D also learns to complete the task although more slowly than PPO.

5.2. The role of ϕ and the value-buffer \mathcal{D}_V

In Figure 5 we mainly compare PPO+D with $\rho = 0.1, \phi = 0.3$ to $\rho = 0.5, \phi = 0.0$. Crucially, the second parameter configuration does not use the buffer \mathcal{D}_V . It is evident in the task "Two boxes easy" that \mathcal{D}_V is essential for solving harder exploration problems. In the "One box easy" task we can see that occasionally on easier tasks not having \mathcal{D}_V can result in faster learning. However, this comes with a greater variance in the training performance of the agent, as sometimes it completely fails to learn.

In Figure 5 we also run vanilla PPO ($\rho = 0.0, \phi = 0.0$) on the task "One box easy" and establish its inability to solve




Figure 5. **Experiments.** Performance for vanilla PPO ($\rho = 0.0, \phi = 0.0$), PPO+D with $\rho = 0.5, \phi = 0.0$ and PPO+D with $\rho = 0.1, \phi = 0.3$ on the tasks "One box easy" and "Two boxes easy" using a single demonstration. Some of the curves overlap each other as they receive zero or close to zero reward. Vanilla PPO never solves the task. The curves show the mean, min and max performance across 5 seeds.

the task even once on any seed and any initial condition. This being the easiest of all tasks, we consider unlikely the event of vanilla PPO successfully completing any of the other harder tasks. We defer an ablation study of the parameter ρ to section B in the Appendix. Although we can not make a direct comparison with (Paine et al., 2019), we think it is useful to underline some of the differences between the two approaches both in the problem setting and in the performance. We attempted to recreate the same complexity of the tasks on which (Paine et al., 2019) was tested. In the article, variability is introduced in the tasks on multiple dimensions such as position and orientation of the agent, shape and color of objects, colors of floor and walls, number of objects in the environment and position of such objects. The range of the initial conditions for each of these factors was not reported. In our work we change the initial position and orientation of the agent as well as the initial positions of the boxes. As for the number of demonstrations in (Paine et al., 2019) the agent has access to a hundred demonstrations, compared to only one in our work. In the present article, the training time ranges from under 5 millions frames to 500 millions, whereas in (Paine et al., 2019) it ranges from 5 billions to 30. Although we adopted the use of the parameter ρ from (Paine et al., 2019) its value differs considerably, which we attribute to the difference between the two algorithms: one being based on PPO, the other on DQN.

It is worth pointing out that PPO+D performance is likely to suffer in environments where the reward function is also stochastic (e.g. imperfect-information games). This is due to the fact that the importance sampling in the current architecture can not compensate for the difference between the training data distribution, which is biased by the replay, and the real environment data distribution.

6. Conclusion

We introduce PPO+D, an algorithm that uses a single demonstration to explore more efficiently in hard-exploration problems. We further improve on the algorithm by introducing two replay buffers: one containing the agent own successful trajectories as it collects these over the training and the second collecting unsuccessful trajectories with a high maximum value estimate. In the second buffer the replay of the trajectories is prioritized according to the maximum estimated value. We show that training with both these buffers solves, to some degree, all tasks the agent was presented with. We also show that vanilla PPO as well as PPO+D without the value-buffer fails to learn the same tasks. In the article, we justify the choice of such adjustments as measures to counter catastrophic forgetting, a problem that afflicts PPO particularly. The present algorithm suffers some limitations as currently it fails to generalize to all variations of some of the problems, yet it achieves to solve several very hard exploration problems with a single demonstration. The source code is available at https://github.com/compsciencelab/ppo_D.

References

- Alex, N. Competing in the obstacle tower challenge. URL <https://blog.aqnichol.com/2019/07/24/competing-in-the-obstacle-tower-challenge/>.
- Badia, A. P., Piot, B., Kapturowski, S., Sprechmann, P., Vitvitskiy, A., Guo, D., and Blundell, C. Agent57: Outperforming the atari human benchmark. *arXiv preprint arXiv:2003.13350*, 2020a.
- Badia, A. P., Sprechmann, P., Vitvitskiy, A., Guo, D., Piot, B., Kapturowski, S., Tieleman, O., Arjovsky, M., Pritzel, A., Bolt, A., et al. Never give up: Learning directed

- exploration strategies. *arXiv preprint arXiv:2002.06038*, 2020b.
- Baker, B., Kanitscheider, I., Markov, T., Wu, Y., Powell, G., McGrew, B., and Mordatch, I. Emergent tool use from multi-agent autotutorials. *arXiv preprint arXiv:1909.07528*, 2019.
- Beattie, C., Leibo, J. Z., Teplyaev, D., Ward, T., Wainwright, M., Küttler, H., Lefrancq, A., Green, S., Valdés, V., Sadik, A., et al. Deepmind lab. *arXiv preprint arXiv:1612.03801*, 2016.
- Bellemare, M., Srinivasan, S., Ostrovski, G., Schaul, T., Saxton, D., and Munos, R. Unifying count-based exploration and intrinsic motivation. In *Advances in Neural Information Processing Systems*, pp. 1471–1479, 2016.
- Bellemare, M. G., Naddaf, Y., Veness, J., and Bowling, M. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.
- Beyret, B., Hernández-Orallo, J., Cheke, L., Halina, M., Shanahan, M., and Crosby, M. The animal-ai environment: Training and testing animal-like artificial cognition. *arXiv preprint arXiv:1909.07483*, 2019.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- Burda, Y., Edwards, H., Storkey, A., and Klimov, O. Exploration by random network distillation. *arXiv preprint arXiv:1810.12894*, 2018.
- Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- Choi, J., Guo, Y., Moczulski, M., Oh, J., Wu, N., Norouzi, M., and Lee, H. Contingency-aware exploration in reinforcement learning. *arXiv preprint arXiv:1811.01483*, 2018.
- Coumans, E. and Bai, Y. Pybullet, a python module for physics simulation in robotics, games and machine learning, 2017.
- Crosby, M., Beyret, B., and Halina, M. The animal-ai olympics. *Nature Machine Intelligence*, 1(5):257, 2019.
- Duan, Y., Andrychowicz, M., Stadie, B. C., Ho, J., Schneider, J., Sutskever, I., Abbeel, P., and Zaremba, W. One-shot imitation learning. *arXiv preprint arXiv:1703.07326*, 2017.
- Ecoffet, A., Huizinga, J., Lehman, J., Stanley, K. O., and Clune, J. Go-explore: a new approach for hard-exploration problems. *arXiv preprint arXiv:1901.10995*, 2019.
- Espeholt, L., Soyer, H., Munos, R., Simonyan, K., Mnih, V., Ward, T., Doron, Y., Firoiu, V., Harley, T., Dunning, I., et al. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. *arXiv preprint arXiv:1802.01561*, 2018.
- Hester, T., Vecerik, M., Pietquin, O., Lanctot, M., Schaul, T., Piot, B., Horgan, D., Quan, J., Sendonaris, A., Osband, I., et al. Deep q-learning from demonstrations. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Ho, J. and Ermon, S. Generative adversarial imitation learning. In *Advances in neural information processing systems*, pp. 4565–4573, 2016.
- Juliani, A., Khalifa, A., Berges, V.-P., Harper, J., Henry, H., Crespi, A., Togelius, J., and Lange, D. Obstacle tower: A generalization challenge in vision, control, and planning. *arXiv preprint arXiv:1902.01378*, 2019.
- Kang, B., Jie, Z., and Feng, J. Policy optimization with demonstrations. In *International Conference on Machine Learning*, pp. 2474–2483, 2018.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kostrikov, I. Pytorch implementations of reinforcement learning algorithms. <https://github.com/ikostrikov/pytorch-a2c-ppo-acktr-gail>, 2018.
- Li, Y., Song, J., and Ermon, S. Infogail: Interpretable imitation learning from visual demonstrations. In *Advances in Neural Information Processing Systems*, pp. 3812–3822, 2017.
- Machado, M. C., Bellemare, M. G., Talvitie, E., Veness, J., Hausknecht, M., and Bowling, M. Revisiting the arcade learning environment: Evaluation protocols and open problems for general agents. *Journal of Artificial Intelligence Research*, 61:523–562, 2018.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529, 2015.
- Oh, J., Guo, Y., Singh, S., and Lee, H. Self-imitation learning. *arXiv preprint arXiv:1806.05635*, 2018.

- Ostrovski, G., Bellemare, M. G., van den Oord, A., and Munos, R. Count-based exploration with neural density models. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 2721–2730. JMLR. org, 2017.
- Paine, T. L., Gulcehre, C., Shahriari, B., Denil, M., Hoffman, M., Soyer, H., Tanburn, R., Kapturowski, S., Rabinowitz, N., Williams, D., et al. Making efficient use of demonstrations to solve hard exploration problems. *arXiv preprint arXiv:1909.01387*, 2019.
- Pathak, D., Agrawal, P., Efros, A. A., and Darrell, T. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 16–17, 2017.
- Rahmatizadeh, R., Abolghasemi, P., Bölöni, L., and Levine, S. Vision-based multi-task manipulation for inexpensive robots using end-to-end learning from demonstration. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3758–3765. IEEE, 2018.
- Ross, S., Gordon, G., and Bagnell, D. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 627–635, 2011.
- Salimans, T. and Chen, R. Learning montezuma’s revenge from a single demonstration. *arXiv preprint arXiv:1812.03381*, 2018.
- Schaul, T., Quan, J., Antonoglou, I., and Silver, D. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*, 2015.
- Schmidhuber, J. Formal theory of creativity, fun, and intrinsic motivation (1990–2010). *IEEE Transactions on Autonomous Mental Development*, 2(3):230–247, 2010.
- Schrittwieser, J., Antonoglou, I., Hubert, T., Simonyan, K., Sifre, L., Schmitt, S., Guez, A., Lockhart, E., Hassabis, D., Graepel, T., et al. Mastering atari, go, chess and shogi by planning with a learned model. *arXiv preprint arXiv:1911.08265*, 2019.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354–359, 2017.
- Srinivas, A., Laskin, M., and Abbeel, P. Curl: Contrastive unsupervised representations for reinforcement learning. *arXiv preprint arXiv:2004.04136*, 2020.
- Tang, H., Houthoofd, R., Foote, D., Stooke, A., Chen, O. X., Duan, Y., Schulman, J., DeTurck, F., and Abbeel, P. Exploration: A study of count-based exploration for deep reinforcement learning. In *Advances in neural information processing systems*, pp. 2753–2762, 2017.
- Williams, R. J. and Peng, J. Function optimization using connectionist reinforcement learning algorithms. *Connection Science*, 3(3):241–268, 1991.
- Zhou, A., Jang, E., Kappler, D., Herzog, A., Khansari, M., Wohlhart, P., Bai, Y., Kalakrishnan, M., Levine, S., and Finn, C. Watch, try, learn: Meta-learning from demonstrations and reward. *arXiv preprint arXiv:1906.03352*, 2019.