



HAL
open science

PBCS: Efficient Exploration and Exploitation Using a Synergy Between Reinforcement Learning and Motion Planning

Guillaume Matheron, Nicolas Perrin, Olivier Sigaud

► **To cite this version:**

Guillaume Matheron, Nicolas Perrin, Olivier Sigaud. PBCS: Efficient Exploration and Exploitation Using a Synergy Between Reinforcement Learning and Motion Planning. Artificial Neural Networks and Machine Learning – ICANN 2020, Sep 2020, Bratislava, Slovakia. pp.295-307, 10.1007/978-3-030-61616-8_24 . hal-03080918

HAL Id: hal-03080918

<https://hal.science/hal-03080918v1>

Submitted on 8 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

PBCS: Efficient Exploration and Exploitation Using a Synergy between Reinforcement Learning and Motion Planning

Guillaume Matheron¹[0000-0001-6530-8784], Nicolas Perrin¹[0000-0003-2358-2915],
and Olivier Sigaud¹[0000-0002-8544-0229]

Sorbonne Université, CNRS, Institut des Systèmes Intelligents et de Robotique, ISIR,
F-75005 Paris, France

Abstract. The exploration-exploitation trade-off is at the heart of reinforcement learning (RL). However, most continuous control benchmarks used in recent RL research only require local exploration. This led to the development of algorithms that have basic exploration capabilities, and behave poorly in benchmarks that require more versatile exploration. For instance, as demonstrated in our empirical study, state-of-the-art RL algorithms such as DDPG and TD3 are unable to steer a point mass in even small 2D mazes. In this paper, we propose a new algorithm called "Plan, Backplay, Chain Skills" (PBCS) that combines motion planning and reinforcement learning to solve hard exploration environments. In a first phase, a motion planning algorithm is used to find a single good trajectory, then an RL algorithm is trained using a curriculum derived from the trajectory, by combining a variant of the Backplay algorithm and skill chaining. We show that this method outperforms state-of-the-art RL algorithms in 2D maze environments of various sizes, and is able to improve on the trajectory obtained by the motion planning phase.

Introduction

Reinforcement Learning (RL) algorithms have been used successfully to optimize policies for both discrete and continuous control problems with high dimensionality [25,23], but fall short when trying to solve difficult exploration problems [17,1,38]. On the other hand, motion planning (MP) algorithms such as RRT [22] are able to efficiently explore in large cluttered environments but, instead of trained policies, they output trajectories that cannot be used directly for closed loop control.

In this paper, we consider environments that present a hard exploration problem with a sparse reward. In this context, a *good trajectory* is one that reaches a state with a positive reward, and we say that an environment is *solved* when a controller is able to reliably reach a rewarded state. We illustrate our approach with 2D continuous action mazes as they facilitate the visual examination of the results, but we believe that this approach can be beneficial to many robotics problems.

If one wants to obtain closed loop controllers for hard exploration problems, a simple approach is to first use an MP algorithm to find a single good trajectory τ , then optimize and robustify it using RL. However, using τ as a stepping stone for an RL

algorithm is not straightforward. In this article, we propose PBCS, an approach that fits the framework of Go-Explore [8], and is based on the Backplay algorithm [35] and skill chaining [20,21]. We show that this approach greatly outperforms both DDPG [23] and TD3 [14] on continuous control problems in 2D mazes, as well as approaches that use Backplay but no skill chaining.

PBCS has two successive phases. First, the environment is explored until a single good trajectory is found. Then this trajectory is used to create a curriculum for training DDPG. More precisely, PBCS progressively increases the difficulty through a backplay process which gradually moves the starting point of the environment backwards along the trajectory resulting from exploration. Unfortunately, this process has its own issues, and DDPG becomes unstable in long training sessions. Calling upon a skill chaining approach, we use the fact that even if Backplay eventually fails, it is still able to solve some subset of the problem. Therefore, a partial policy is saved, and the remainder of the problem is solved recursively until the full environment can be solved reliably.

In this article, we contribute an extension of the Go-Explore framework to continuous control environments, a new way to combine a variant of the Backplay algorithm with skill chaining, and a new state-space exploration algorithm.

1 Related Work

Many works have tried to incorporate better exploration mechanisms in RL, with various approaches.

Encouraging exploration of novel states The authors of [43] use a count-based method to penalize states that have already been visited, while the method proposed by [2] reuses actions that have provided diverse results in the past. Some methods try to choose policies that are both efficient and novel [33,7,34,9], while some use novelty as the only target, entirely removing the need for rewards [10,19]. The authors of [41,3,31] train a forward model and use the unexpectedness of the environment step as a proxy for novelty, which is encouraged through reward shaping. Some approaches try to either estimate the uncertainty of value estimates [29], or learn bounds on the value function [5]. All these solutions try to integrate an exploration component within RL algorithms, while our approach separates exploration and exploitation into two successive phases, as in [6].

Using additional information about the environment Usually in RL, the agent can only learn about the environment through interactions. However, when additional information about the task at hand is provided, other methods are available. This information can take the form of expert demonstrations [37,18,35,21,13,27,30], or having access to a single rewarded state [12]. When a full representation of the environment is known, RL can still be valuable to handle the dynamics of the problem: PRM-RL [11] and RL-RRT [4] use RL as reachability estimators during a motion planning process.

Building on the Go-Explore framework To our knowledge, the closest approach to ours is the Go-Explore [8] framework, but in contrast to PBCS, Go-Explore is applied to

discrete problems such as Atari benchmarks. In a first phase, a single valid trajectory is computed using an ad-hoc exploration algorithm. In a second phase, a learning from demonstration (LfD) algorithm is used to imitate and improve upon this trajectory. Go-Explore uses Backplay [35,37] as the LfD algorithm, with Proximal Policy Optimization (PPO) [40] as policy optimization method. Similar to Backplay, the authors of [16] have proposed Recall Traces, a process in which a backtracking model is used to generate a collection of trajectories reaching the goal.

The authors of [26] present an approach that is similar to ours, and also fits the framework of Go-Explore. In phase 1, they use a guided variant of RRT, and in phase 2 they use a learning from demonstration algorithm based on TRPO. Similarly, PBCS follows the same two phases as Go-Explore, with major changes to both phases. In the first phase, our exploration process is adapted to continuous control environments by using a different binning method, and different criteria for choosing the state to reset to. In the second phase, a variant of Backplay is integrated with DDPG instead of PPO, and seamlessly integrated with a skill chaining strategy and reward shaping.

The Backplay algorithm in PBCS is a deterministic variant of the one proposed in [35]. In the original Backplay algorithm, the starting point of each policy is chosen randomly from a subset of the trajectory, but in our variant the starting point is deterministic: the last state of the trajectory is used until the performance of DDPG converges (more details are presented in Sect. 8), then the previous state is chosen, and so on until the full trajectory has been exploited.

Skill chaining The process of *skill chaining* was explored in different contexts by several research papers. The authors of [20] present an algorithm that incrementally learns a set of skills using classifiers to identify changepoints, while the method proposed in [21] builds a skill tree from demonstration trajectories, and automatically detects changepoints using statistics on the value function. To our knowledge, our approach is the first to use Backplay to build a skill chain. We believe that it is more reliable and minimizes the number of changepoints because the position of changepoints is decided using data from the RL algorithm that trains the policies involved in each skill.

2 Background

Our work is an extension of the Go-Explore algorithm. In this section, we summarize the main concepts of our approach.

Reset-anywhere. Our work makes heavy use of the ability to reset an environment to any state. The use of this primitive is relatively uncommon in RL, because it is not always readily available, especially in real-world robotics problems. However, it can be invaluable to speed up exploration of large state spaces. It was used in the context of Atari games by [18], proposed in [39] as VINE, and gained popularity with [37].

Sparse rewards. Most traditional RL benchmarks only require very local exploration, and have smooth rewards guiding them towards the right behavior. Thus sparse rewards problems are especially hard for RL algorithms: the agent has to discover without any

external signal a long sequence of actions leading to the reward. Most methods that have been used to help with this issue require prior environment-specific knowledge [36].

Maze environments. Lower-dimension environments such as cliff walk [42] are often used to demonstrate fundamental properties of RL algorithms, and testing in these environments occasionally reveals fundamental flaws [24]. We deliberately chose to test our approach on 2D maze environments because they are hard exploration problems, and because reward shaping behaves very poorly in such environments, creating many local optima. Our results in Section 4 show that state-of-the-art algorithms such as DDPG and TD3 fail to solve even very simple mazes.

DDPG. Deep Deterministic Policy Gradient (DDPG) is a continuous action actor-critic algorithm using a deterministic actor that performs well on many control tasks [23]. However, DDPG suffers from several sources of instability. Our maze environments fit the analysis made by the authors of [32], according to whom the critic approximator may "leak" Q-value across walls of discontinuous environments. With a slightly different approach, [15] suggests that extrapolation error may cause DDPG to over-estimate the value of states that have never been visited or are unreachable, causing instability. More generally, the authors of [42] formalize the concept of "deadly triad", according to which algorithms that combine function approximation, bootstrapping updates and off-policy are prone to diverge. Even if the deadly triad is generally studied in the context of the DQN algorithm [25], these studies could also apply to DDPG. Finally, the authors of [24] show that DDPG can fail even in trivial environments, when the reward is not found quickly enough by the built-in exploration of DDPG.

3 Methods

Figure 1 describes PBCS. The algorithm is split in two successive phases, mirroring the Go-Explore framework. In a first phase, the environment is incrementally explored until a single rewarded state is found. In a second phase, a single trajectory provides a list of starting points, that are used to train DDPG on increasingly difficult portions of the full environment. Each time the problem becomes too difficult and DDPG starts to fail, training stops, and the trained agent is recorded as a local skill. Training then resumes for the next skill, with a new target state. This loop generates a set of skills that can then be chained together to create a controller that reaches the target reliably.

Notations

State neighborhood. For any state $\tau_i \in S$, and $\epsilon > 0$, we define $B_\epsilon(\tau_i)$ as the closed ball of radius ϵ centered around τ_i . Formally, this is the set $\{s \in S \mid d(s, \tau_i) \leq \epsilon\}$ where d is the L2 distance.

Skill chaining. Skill chaining consists in splitting a complex task into simpler sub-tasks that are each governed by a different policy. Complex tasks can then be solved by executing each policy sequentially.

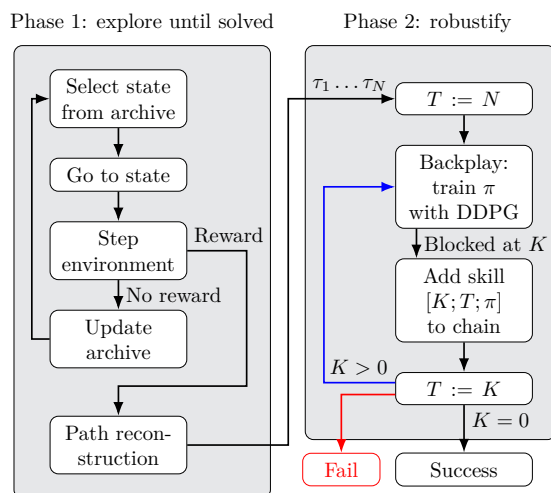


Fig. 1: Overview of PBCS. The red path is only used when testing the algorithm without skill chaining, otherwise the blue path is used.

Formally, each task T_i has an activation condition $A_i \subset S$, and a policy $\pi_i : S \rightarrow A$. A *task chain* is a list of tasks $T_0 \dots T_n$, which can be executed sequentially: the actor uses π_0 until the state of the system reaches a state $s \in A_1$, then it uses π_1 , and so on until the end of the episode (which can be triggered by reaching either a terminal state or a predetermined maximum number of steps).

3.1 Phase 1: Explore until Solved

In phase 1, PBCS explores to find a single path that obtains a non-zero reward in the environment. This exploration phase is summarized in this section, and implementation details are available in Appendix S1. An archive keeps track of all the visited states. In this archive, states s are grouped in square state-space bins. A state-counter c_s is attached to each state, and a bin-counter c_b is attached to each bin. All counters are initialized to 0.

The algorithm proceeds in 5 steps, as depicted in Fig. 1:

1. **Select state from archive.** To select a state, the non-empty bin with the lowest counter is first selected, then from all the states in this bin, the state with the lowest counter is selected. Both the bin and state counters are then incremented.
2. **Go to state.** The environment is reset to the selected state. This assumes the existence of a "reset-anywhere" primitive, which can be made available in simulated environments.
3. **Step environment.** A single environment step is performed, with a random action.
4. **Update archive.** The newly-reached state is added to the archive if not already present.

5. **Termination of phase 1.** As soon as the reward is reached, the archive is used to reconstruct the sequence of states that led the agent from its initial state to the reward. This sequence $\tau_0 \dots \tau_N$ is passed on to phase 2.

This process can be seen as a random walk with a constraint on the maximum distance between two states: in the beginning, a single trajectory is explored until it reaches a state that is too close to an already-visited state. When this happens, a random visited state is selected as the starting point of a new random walk. Another interpretation of this process is the construction of a set of states with uniform spatial density. Under this view, the number of states in each cell is used as a proxy for the spatial density of the distribution.

3.2 Phase 2: Robustify

Phase 2 of PBCS learns a controller from the trajectory obtained in phase 1.

Algorithm 1: Phase 2 of PBCS

Input : $\tau_0 \dots \tau_N$ the output of phase 1
Output: $\pi_0 \dots \pi_n$ a chain of policies with activation sets $A_0 \dots A_n$

- 1 $T = N$
- 2 $n = 0$
- 3 **while** $T > 0$ **do**
- 4 $\pi_n, T = \text{Backplay}(\tau_0 \dots \tau_T)$
- 5 $A_n = B_\epsilon(\tau_T)$
- 6 $n = n + 1$
- 7 **end**
- 8 Reverse lists $\pi_0 \dots \pi_n$ and $A_0 \dots A_n$

Skill Chaining. Algorithm 1 presents the skill chaining process. It uses the `Backplay` function, that takes as input a trajectory $\tau_0 \dots \tau_T$, and returns a policy π and an index $K < T$ such that running policy π repeatedly on a state from $B_\epsilon(\tau_K)$ always leads to a state in $B_\epsilon(\tau_T)$. The main loop builds a chain of skills that roughly follows trajectory τ , but is able to improve upon it. Specifically, activation sets A_n are centered around points of τ but policies π_n are constructed using a generic RL algorithm that optimizes the path between two activation sets. The list of skills is then reversed, because it was constructed backwards.

Backplay. The `Backplay` algorithm was originally proposed in [35]. More details on the differences between this original algorithm and our variant are available in sections 1 and S6.

The `Backplay` function (Algorithm 2) takes as input a section $\tau_0 \dots \tau_T$ of the trajectory obtained in phase 1, and returns a (K, π) pair where K is an index on trajectory τ , and π is a policy trained to reliably attain $B_\epsilon(\tau_T)$ from $B_\epsilon(\tau_K)$. The policy

π is trained using DDPG to reach $B_\epsilon(\tau_T)$ from starting point $B_\epsilon(\tau_K)$ ¹, where K is initialized to $T - 1$, and gradually decremented in the main loop.

At each iteration, the algorithm evaluates the feasibility of a skill with target $B_\epsilon(\tau_T)$, policy π and activation set $B_\epsilon(\tau_K)$. If the measured performance is 100% without any training (line 5), the current skill is saved and the starting point is decremented. Otherwise, a training loop is executed until performance stabilizes (line 8). This is performed by running Algorithm 3 repeatedly until no improvement over the maximum performance is observed α times in a row. We ran our experiments with $\alpha = 10$, and a more in-depth discussion of hyperparameters is available in Appendix S2.

Then the performance of the skill is measured again (line 9), and three cases are handled:

- **The skill is always successful (line 10).** The current skill is saved and the index of the starting point is decremented.
- **The skill is never successful (line 13).** The last successful skill is returned.
- **The skill is sometimes successful.** The current skill is not saved, and the index of the starting point is decremented. In our maze environment, this happens when $B_\epsilon(\tau_K)$ overlaps a wall: in this case some states of $B_\epsilon(\tau_K)$ cannot reach the target no matter the policy.

Reward Shaping. With reward shaping, we bypass the reward function of the environment, and train DDPG to reach any state τ_T . We chose to use the method proposed by [28]: we define a potential function in Equation (1a), where $d(s, A_i)$ is the L2 distance between s and the center of A_i . We then define our shaped reward in Equation (1b).

$$\Phi(s) = \frac{1}{d(s, A_i)} \tag{1a}$$

$$R_{\text{shaped}}(s, a, s') = \begin{cases} 10 & \text{if } s \in A_i \\ \Phi(s') - \Phi(s) & \text{otherwise.} \end{cases} \tag{1b}$$

Algorithm 3 shows how this reward function is used in place of the environment reward. This training function runs β episodes of up to `max_steps` steps each, and returns the fraction of episodes that were able to reach the reward. β is a hyperparameter that we set to 50 for our test, and more details on this choice are available in Appendix S2.

Importantly, reaching a performance of 100% is not always possible, even with long training sessions, because the starting point is selected in $B_\epsilon(\tau_K)$, and some of these states may be inside obstacles for instance.

Algorithm 2: The Backplay algorithm

Input : $(\tau_0 \dots \tau_T)$ a state-space trajectory
Output: π_s a trained policy
 K_s the index of the starting point of the policy





















```
1  $K = T - 1$ 
2 Initialize a DDPG architecture with policy  $\pi$ 
3 while  $K > 0$  do
4   Test performance of  $\pi$  between  $B_\epsilon(\tau_K)$  and  $B_\epsilon(\tau_T)$  over  $\beta$  episodes
5   if performance = 100% then
6     |  $\pi_s = \pi, K_s = K$ 
7   else
8     | Run Train (Algorithm 3) repeatedly until performance stabilizes.
9     | Test performance of  $\pi$  between  $B_\epsilon(\tau_K)$  and  $B_\epsilon(\tau_T)$  over  $\beta$  episodes
10    | if performance = 100% then
11      |  $\pi_s = \pi, K_s = K$ 
12    | end
13    | if performance = 0% and  $K_s$  exists then
14      | return  $(K_s, \pi_s)$ 
15    | end
16  end
17   $K = K - 1$ 
18 end
19 return  $(K_s, \pi_s)$ 
```

Algorithm 3: Training process with reward shaping

Input : τ_K the source state
 τ_T the target state
Output: The average performance p

```
1  $n = 0$ 
2 for  $i = 1 \dots \beta$  do
3   |  $s \sim B_\epsilon(\tau_K)$ 
4   | for  $j = 1 \dots max\_steps$  do
5     |  $a = \pi(s) + \text{random noise}$ 
6     |  $s' = \text{step}(s, a)$ 
7     |  $r = \begin{cases} 10 & d(s', \tau_T) \leq \epsilon \\ \frac{1}{d(s', \tau_T)} - \frac{1}{d(s, \tau_T)} & \text{otherwise} \end{cases}$ 
8     | DDPG.train( $s, a, s', r$ )
9     |  $s = s'$ 
10    | if  $d(s', \tau_T) \leq \epsilon$  then
11      |  $n = n + 1$ 
12    | break
13  | end
14  end
15 end
16  $p = \frac{n}{\beta}$ 
```

Table 1: Results of various algorithms on maze environments. For each test, the number of environment steps performed is displayed with a red background when the policy was not able to reach the target, and a green one when training was successful. In "Vanilla" experiments, the red paths represent the whole area explored by the RL algorithm. In "Backplay" experiments, the trajectory computed in phase 1 is displayed in red, and the "robustified" policy or policy chain is displayed in green. Activation sets A_i are displayed as purple circles. Enlarged images are presented in Fig. S2.

Vanilla		PBCS w/o skill chaining	PBCS
DDPG	TD3	DDPG	DDPG
 1M	 1M	 146k	 321k
 1M	 1M	 372k	 5M
 1M	 1M	 268k	 6M
 1M	 1M	 694k	 8M
 1M	 1M	 175k	 22M

4 Experimental Results

We perform experiments in continuous maze environments of various sizes. For a maze of size N , the state-space is the position of a point mass in $[0, N]^2$ and the action describes the speed of the point mass, in $[-0.1, 0.1]^2$. Therefore, the step function is simply $s' = s + a$, unless the $[s, s']$ segment intersects a wall. The only reward is -1 when hitting a wall and 1 when the target area is reached. A more formal definition of the environment is available in Appendix S4.

Our results are presented in Table 1. We first tested standard RL algorithms (DDPG and TD3), then PBCS, but without skill chaining (this was done by replacing the blue branch with the red branch in Fig. 1). When the full algorithm would add a new skill to the skill chain and continue training, this variant stops and fails. These results are presented in column "PBCS without skill chaining". Finally, the full version of PBCS with skill chaining is able to solve complex mazes up to 15×15 cells, by chaining several intermediate skills.

5 Discussion of Results

As expected, standard RL algorithms (DDPG and TD3) were unable to solve all but the simplest mazes. These algorithms have no mechanism for state-space exploration other than uniform noise added to their policies during rollouts. Therefore, in the best-case scenario they perform a random walk and, in the worst-case scenario, their actors may actively hinder exploration.

More surprisingly, PBCS without skill chaining is still unable to reliably² solve mazes larger than 2×2 . Although phase 1 always succeeds in finding a feasible trajectory τ , the robustification phase fails relatively early. We attribute these failures to well-known limitations of DDPG exposed in Section 2. We found that the success rate of PBCS without skill chaining was very dependent on the discount rate γ , which we discuss in Appendix S3.

The full version of PBCS with skill chaining is able to overcome these issues by limiting the length of training sessions of DDPG, and is able to solve complex mazes up to 7×7 , by chaining several intermediate skills.

6 Conclusion

The authors of Go-Explore identified state-space exploration as a fundamental difficulty on two Atari benchmarks. We believe that this difficulty is also present in many continuous control problems, especially in high-dimension environments. We have shown that the PBCS algorithm can solve these hard exploration, continuous control environments

¹ More details on why the starting point needs to be $B_\epsilon(\tau_K)$ instead of τ_K are available in Appendix S6

² We tested PBCS without skill chaining with different seeds on small mazes, these results are presented in Appendix S3

by combining a motion planning process with reinforcement learning and skill chaining. Further developments should focus on testing these hybrid approaches on higher dimensional environments that present difficult exploration challenges together with difficult local control, such as the Ant-Maze MuJoCo benchmark [44], and developing methods that use heuristics suited to continuous control in the exploration process, such as Quality-Diversity approaches [34].

7 Acknowledgements

This work was partially supported by the French National Research Agency (ANR), Project ANR-18-CE33-0005 HUSKI.

References

1. Achiam, J., Knight, E., Abbeel, P.: Towards Characterizing Divergence in Deep Q-Learning. arXiv:1903.08894 (2019)
2. Benureau, F.C.Y., Oudeyer, P.Y.: Behavioral Diversity Generation in Autonomous Exploration through Reuse of Past Experience. *Front. Robot. AI* **3** (2016)
3. Burda, Y., Edwards, H., Storkey, A., Klimov, O.: Exploration by Random Network Distillation. arXiv:1810.12894 (2018)
4. Chiang, H.T.L., Hsu, J., Fiser, M., Tapia, L., Faust, A.: RL-RRT: Kinodynamic Motion Planning via Learning Reachability Estimators from RL Policies. arXiv:1907.04799 (2019)
5. Ciosek, K., Vuong, Q., Loftin, R., Hofmann, K.: Better Exploration with Optimistic Actor-Critic. arXiv:1910.12807 (2019)
6. Colas, C., Sigaud, O., Oudeyer, P.Y.: GEP-PG: Decoupling Exploration and Exploitation in Deep Reinforcement Learning Algorithms. arXiv:1802.05054 (2018)
7. Cully, A., Demiris, Y.: Quality and Diversity Optimization: A Unifying Modular Framework. *IEEE Transactions on Evolutionary Computation* pp. 1–1 (2017)
8. Ecoffet, A., Huizinga, J., Lehman, J., Stanley, K.O., Clune, J.: Go-Explore: a New Approach for Hard-Exploration Problems. arXiv:1901.10995 (2019)
9. Erickson, L.H., LaValle, S.M.: Survivability: Measuring and ensuring path diversity. In: 2009 IEEE International Conference on Robotics and Automation. pp. 2068–2073 (2009)
10. Eysenbach, B., Gupta, A., Ibarz, J., Levine, S.: Diversity is All You Need: Learning Skills without a Reward Function. arXiv:1802.06070 (2018)
11. Faust, A., Ramirez, O., Fiser, M., Oslund, K., Francis, A., Davidson, J., Tapia, L.: PRM-RL: Long-range Robotic Navigation Tasks by Combining Reinforcement Learning and Sampling-based Planning. arXiv:1710.03937 (2018)
12. Florensa, C., Held, D., Wulfmeier, M., Zhang, M., Abbeel, P.: Reverse Curriculum Generation for Reinforcement Learning. arXiv:1707.05300 (2018)
13. Fournier, P., Sigaud, O., Colas, C., Chetouani, M.: CLIC: Curriculum Learning and Imitation for object Control in non-rewarding environments. arXiv:1901.09720 (2019)
14. Fujimoto, S., Hoof, H.v., Meger, D.: Addressing Function Approximation Error in Actor-Critic Methods. *ICML* (2018)
15. Fujimoto, S., Meger, D., Precup, D.: Off-Policy Deep Reinforcement Learning without Exploration. arXiv:1812.02900 (2018)
16. Goyal, A., Brakel, P., Fedus, W., Singhal, S., Lillicrap, T., Levine, S., Larochelle, H., Bengio, Y.: Recall Traces: Backtracking Models for Efficient Reinforcement Learning. arXiv:1804.00379 (2019)

17. van Hasselt, H., Doron, Y., Strub, F., Hessel, M., Sonnerat, N., Modayil, J.: Deep Reinforcement Learning and the Deadly Triad. arXiv:1812.02648 (2018)
18. Hosu, I.A., Rebedea, T.: Playing Atari Games with Deep Reinforcement Learning and Human Checkpoint Replay. arXiv:1607.05077 (2016)
19. Knepper, R.A., Mason, M.T.: Path diversity is only part of the problem. In: 2009 IEEE International Conference on Robotics and Automation. pp. 3224–3229 (2009)
20. Konidaris, G., Barto, A.G.: Skill Discovery in Continuous Reinforcement Learning Domains using Skill Chaining. In: Bengio, Y., et. al. (eds.) Advances in Neural Information Processing Systems 22, pp. 1015–1023 (2009)
21. Konidaris, G., Kuindersma, S., Grupen, R., Barto, A.G.: Constructing Skill Trees for Reinforcement Learning Agents from Demonstration Trajectories. In: Lafferty, J.D., et. al. (eds.) Advances in Neural Information Processing Systems 23, pp. 1162–1170 (2010)
22. Lavalle, S.M.: Rapidly-Exploring Random Trees: A New Tool for Path Planning. Tech. rep., Iowa State University (1998)
23. Lillicrap, T.P., Hunt, J.J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., Wierstra, D.: Continuous control with deep reinforcement learning. arXiv:1509.02971 (2015)
24. Matheron, G., Perrin, N., Sigaud, O.: The problem with DDPG: understanding failures in deterministic environments with sparse rewards. arXiv:1911.11679 (2019)
25. Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., Riedmiller, M.: Playing Atari with Deep Reinforcement Learning. arXiv:1312.5602 (2013)
26. Morere, P., Francis, G., Blau, T., Ramos, F.: Reinforcement Learning with Probabilistically Complete Exploration. arXiv:2001.06940 (2020)
27. Nair, A., McGrew, B., Andrychowicz, M., Zaremba, W., Abbeel, P.: Overcoming Exploration in Reinforcement Learning with Demonstrations. arXiv:1709.10089 (2018)
28. Ng, A.Y., Harada, D., Russell, S.J.: Policy Invariance Under Reward Transformations: Theory and Application to Reward Shaping. In: Proceedings of the Sixteenth International Conference on Machine Learning. pp. 278–287. ICML '99 (1999)
29. Osband, I., Blundell, C., Pritzel, A., Van Roy, B.: Deep Exploration via Bootstrapped DQN. arXiv:1602.04621 (2016)
30. Paine, T.L., Gulcehre, C., Shahriari, B., Denil, M., Hoffman, M., Soyer, H., Tanburn, R., Kaputowski, S., Rabinowitz, N., Williams, D., Barth-Maron, G., Wang, Z., de Freitas, N.: Making Efficient Use of Demonstrations to Solve Hard Exploration Problems. arXiv:1909.01387 (2019)
31. Pathak, D., Agrawal, P., Efros, A.A., Darrell, T.: Curiosity-driven Exploration by Self-supervised Prediction. arXiv:1705.05363 (2017)
32. Penedones, H., Vincent, D., Maennel, H., Gelly, S., Mann, T., Barreto, A.: Temporal Difference Learning with Neural Networks - Study of the Leakage Propagation Problem. arXiv:1807.03064 (2018)
33. Pugh, J.K., Soros, L.B., Szerlip, P.A., Stanley, K.O.: Confronting the Challenge of Quality Diversity. In: Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation. pp. 967–974. GECCO '15, ACM, New York, NY, USA (2015)
34. Pugh, J.K., Soros, L.B., Stanley, K.O.: Quality Diversity: A New Frontier for Evolutionary Computation. *Front. Robot. AI* **3** (2016)
35. Resnick, C., Raileanu, R., Kapoor, S., Peysakhovich, A., Cho, K., Bruna, J.: Backplay: "Man muss immer umkehren". arXiv:1807.06919 (2018)
36. Riedmiller, M., Hafner, R., Lampe, T., Neunert, M., Degraeve, J., Van de Wiele, T., Mnih, V., Heess, N., Springenberg, J.T.: Learning by Playing - Solving Sparse Reward Tasks from Scratch. arXiv:1802.10567 (2018)
37. Salimans, T., Chen, R.: Learning Montezuma's Revenge from a Single Demonstration. arXiv:1812.03381 (2018)

38. Schaul, T., Quan, J., Antonoglou, I., Silver, D.: Prioritized Experience Replay. arXiv:1511.05952 (2015)
39. Schulman, J., Levine, S., Moritz, P., Jordan, M.I., Abbeel, P.: Trust Region Policy Optimization. arXiv:1502.05477 (2015)
40. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal Policy Optimization Algorithms. arXiv:1707.06347 (2017)
41. Stadie, B.C., Levine, S., Abbeel, P.: Incentivizing Exploration In Reinforcement Learning With Deep Predictive Models. arXiv:1507.00814 (2015)
42. Sutton, R.S., Barto, A.G.: Reinforcement Learning: An Introduction. MIT Press (2018)
43. Tang, H., Houthoofd, R., Foote, D., Stooke, A., Chen, X., Duan, Y., Schulman, J., De Turck, F., Abbeel, P.: #Exploration: A Study of Count-Based Exploration for Deep Reinforcement Learning. arXiv:1611.04717 (2016)
44. Tassa, Y., et. al.: DeepMind Control Suite. arXiv:1801.00690 (2018)
45. Wilson, E.B.: Probable Inference, the Law of Succession, and Statistical Inference. *Journal of the American Statistical Association* **22**(158), 209–212 (1927)

Supplemental Materials

S1 Phase 1: Explore Until Solved

Algorithm 4: Exploration algorithm

Input : $s_0 \in S$ the initial environment state
step : $S \times A \rightarrow S \times \mathbb{R} \times \mathbb{B}$ the environment step function
iterations $\in \mathbb{N}$ the number of samples to accumulate
Bin : $F \rightarrow \mathbb{N}$ a binning function
Output: transitions $\subseteq S \times A \times \mathbb{R} \times \mathbb{B} \times S$ the set of explored transitions

```
1 transitions =  $\emptyset$ 
2  $b_0 = \text{Bin}(s_0)$ 
3 bin_usage[ $b_0$ ] = 0
4 states_in_bin[ $b_0$ ] =  $\{s_0\}$ 
5 state_usage[ $s_0$ ] = 0
6  $B = B \cup \{b_0\}$ 
7 while |transitions| < iterations do
8     chosen_bin = argmin $b \in B$  bin_usage[b]
9     chosen_state = argmin $s \in \text{states\_in\_bin}[\text{chosen\_bin}]$  state_usage[chosen_state]
10    action = random_action()
11    s',reward,terminal = step(chosen_state, action)
12    bin_usage[chosen_bin] ++
13    state_usage[chosen_state] ++
14    transitions = transitions  $\cup$  (chosen_state,action,reward,terminal,s')
15    if NOT terminal then
16        state_usage[s'] = 0
17        b' = Bin(s')
18        if  $b' \in B$  then
19            states_in_bin[b'] = states_in_bin[b']  $\cup$  {s'}
20        else
21            B = B  $\cup$  {b'}
22            bin_usage[b'] = 0
23            states_in_bin[b'] = {s'}
24        end
25    end
26 end
```

Our proposed phase 1 exploration algorithm maintains a pool of states, which initially contains only the start state. At each step, a selection process described below is used to select a state s from the pool (lines 8 to 9 in Algorithm 4). A random action a is then chosen, and the environment is used to compute a single step from state s using

action a (line 11). If the resulting state s' is non-terminal, then it is added to the pool (lines 16 to 23). This process is repeated as long as necessary.

State selection The pool of states stored by the algorithm is divided into square bins of size 0.05. During the state selection process (lines 8 to 9), the least-chosen bin is selected (on line 8), then the least-chosen state from this bin is selected (line 9). In both cases, when several states or bins are tied in the argmin operation, one is selected uniformly randomly from the set of all tied elements.

S2 Choice of PBCS Hyperparameters

PBCS uses three hyperparameters α , β and ϵ .

The parameter α represents the number of consecutive non-improvements of the training performance required to assume training is finished. In our experiments, this value was set to 10, and we summarize here what can be expected if this parameter is set too high or too low.

- Setting α too high results in longer training sessions, in which the policy keeps being trained despite being already successful. The time and sample performance of PBCS is impacted, but the algorithm should still be able to build a policy chain.
- Setting α too low may cause training to stop early. In benign cases, the policy is simply sub-optimal, but in some cases this may lead to the creation of many changepoints, and prevent PBCS from improving at all upon the phase 1 trajectory. If the activation conditions overlap too much, PBCS may output a skill chain that is unable to navigate the environment.

The parameter β represents the number of samples used to evaluate the performance of a skill. In our experiments, this value was set to 50.

- Setting β too high would increase the time and sample complexity of PBCS, but would not impact the output.
- The main risk of setting β too low is that PBCS may incorrectly compute that a skill has a performance of 100%. If this skill is then selected, the output skill chain may be unable to navigate the environment.

The parameter ϵ corresponds to the radius of the targets used during skill chaining.

S3 Choice of Discount Factor γ

The discount factor γ is usually considered to be a parameter of the environment and not the RL algorithm. It controls the decay that is applied when evaluated the contributions of future rewards to a present choice. In our experiments, we tested two values of γ , that are $\gamma = 0.9$ and $\gamma = 0.99$.

DDPG uses a neural network \hat{Q} in order to estimate the state-action value function $Q^\pi(s, a)$ of the current policy π . In the case of deterministic environments, the state-action value function is recursively defined as $Q^\pi(s, a) = R(s, a, s') + \gamma Q(s', \pi(a))$, where $s' = \text{step}(s, a)$.

Therefore, reaching a sparse reward of value 1 after n steps with no reward carries a discounted value of γ^n . This implies that rewards that are reached only after many steps have very little impact on the shape of Q . For instance, with $\gamma = 0.9$ and $n = 50$, $\gamma^n \approx 0.05$. This effectively reduces the magnitude of the training signal used by the actor update of DDPG, reducing the speed of actor updates the farther from the reward. Evidence of this is presented in Sect. S3.1.

With this consideration, it seems that choosing γ very close to 1 solves the problem of exponential decay of the training signal. However, high γ values present their own challenges. The state-action critic approximator \hat{Q} used by DDPG is trained on (s, a, r, s') tuples stored in an experience replay buffer, but as with any continuous approximator, it generalizes the training data to nearby (s, a) couples.

In environments with positive rewards, \hat{Q} can over-estimate the value of states: for instance in maze environments, the learned value can be generalized incorrectly and "leak" through walls.

This mechanism is usually counter-balanced by the fact that over-estimated $Q(s, a)$ values can then be lowered. For instance, in our maze environments, hitting a wall generates a training tuple with $s' = s$ and $r = 0$. The update rule of DDPG applied to this tuple yields: $Q(s, a) \leftarrow Q(s, a)(1 + c(\gamma - 1))$ where c is the critic learning rate. Therefore, the closer γ is to 1, the slower over-estimated values will be corrected.

In smaller mazes, our experiments show that reducing gamma increases the performance of PBCS without skill chaining (Fig. S1).

S3.1 Replacing the Actor Update of DDPG

We claim that the lower reward signal obtained with γ values close to 1 affect the actor update of DDPG. We can test this claim by using a variation of DDPG proposed by the authors of [24]: we replace the actor update of DDPG with a brute-force approach that, for each sampled state s , computes $\max_a \hat{Q}(s, a)$ using uniform sampling. The performance of this variant is presented in Fig. S1 with green and red bars.

S4 Experimental Setup

Our experiments are conducted in maze environments of various sizes. A maze of size N is described using the following Markov Decision Process:

$$\begin{aligned}
 S &= [0, N] \times [0, N] \\
 A &= [-0.1, 0.1] \times [-0.1, 0.1] \\
 R(s, a, s') &= \mathbb{1}_{\|s' - \text{target}\| < 0.2} - \mathbb{1}_{[s, s'] \text{ intersects a wall}} \\
 \text{step}(s, a) &= \begin{cases} s & \text{if } [s, s + a] \text{ intersects a wall} \\ s + a & \text{otherwise.} \end{cases}
 \end{aligned}$$

The set of walls is constructed using a maze generation algorithm, and walls have a thickness of 0.1.

The target position is $(N - .5, N - .5)$ when $N > 2$. In mazes of size 2, the target position is $(.5, 1.5)$.




Fig. S1: Success rate of PBCS without skill chaining, depending on γ . Bars marked with *(am)* use the variant of DDPG presented in Sect. S3.1. Error bars are computed using Wilson score intervals [45].

S5 Enlarged Results

A more detailed view of the results of PBCS on mazes of different sizes is presented in Fig. 1.

S6 Need for Resetting in Unseen States

As a reminder, for the Backplay algorithm and our variant, a single trajectory $\tau_0 \dots \tau_T$ is provided, and training is performed by changing the starting point of the environment to various states.

In the original Backplay algorithm, the environment is always reset to a visited state τ_K , where K is an index chosen randomly in a sliding window of $[0, T]$. The sliding window is controlled by hyperparameters, but the main idea is that in the early stages of training, states near T are more likely to be selected, and in later stages, states near 0 are more likely to be used.

However, we found that this caused a major issue when combined with continuous control and the skill chaining process. With skill chaining, the algorithm creates a sequence of activation sets (A_n) , and a sequence of policies (π_n) such that when the agent reaches a state in A_n , it switches to policy π_n . Each activation set A_n is a ball of radius ϵ centered around a state τ_K for some K .


The policy needs to be trained not only on portions of the environment that are increasingly long, it also needs to account for the uncertainty of its starting point. When executing the skill chain, the controller switched to policy π_n as soon as the state reaches the activation set A_n , which is $B_\epsilon(\tau_K)$ for some K . Even if A_n is relatively small, we found it caused systematic issues on maze environments, as presented in Fig. S3.

In our variant of the Backplay algorithm, we found it was necessary to train DDPG on starting points chosen randomly in $B_\epsilon(\tau_K)$, to ensure that the policy is trained correctly to solve a portion of the environment with any starting point in this volume.

This also means that we need to reset the environment to unseen states, and can cause problems when these states are unreachable (in our maze examples this is usually because they are inside walls, but in higher dimensions we assume this could be more problematic).

When possible, a solution would be to run the environment backwards from τ_K with random actions to generate these samples (while ensuring that they still lie within $B_\epsilon(\tau_K)$). Another solution, especially in high-dimension environments, would be to run the environment backwards for a fixed number of steps, and use a classifier to define the bounds of A_n , instead of using the L2 distance.

Fig. S2: Enlarged view of the results of PBCS on mazes of different sizes. The trajectory computed in phase 1 is displayed in red, and the "robustified" policy chain is displayed in green. Activation sets A_i are displayed as purple circles.



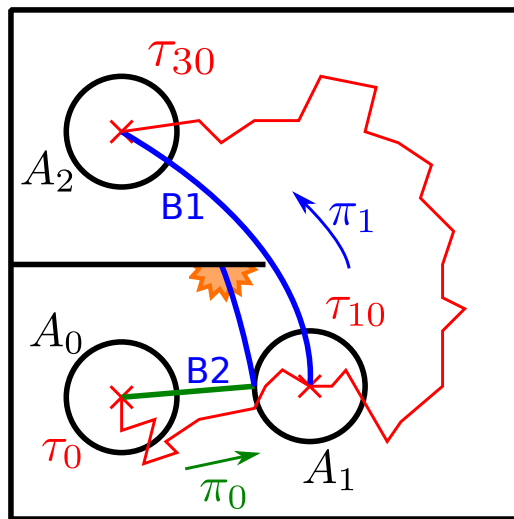


Fig. S3: Policy π_1 was trained using starting points $\tau_{30} \dots \tau_{10}$ without any added noise. Therefore, τ_{30} is reachable from τ_{10} using π_1 (trajectory $B2$), but not necessarily from any point in A_1 . In maze environments, the optimal policy is usually close to walls, and provides little margin for perturbations. The trajectory $B2$ (that starts in green and ends in blue) results from the execution of the skill chain. The controller switches from π_0 to π_1 as soon as the agent reaches A_1 , and then hits the wall (trajectory $B2$). This problem persists even when ϵ is reduced.