

---

# Exploring Restart Distributions

---

Arash Tavakoli<sup>1\*</sup>

Vitaly Levdik<sup>1\*</sup>

Riashat Islam<sup>2,3</sup>

Christopher M. Smith<sup>1</sup>

Petar Kormushev<sup>1</sup>

<sup>1</sup>Imperial College London, <sup>2</sup>Mila, <sup>3</sup>McGill University

## Abstract

We consider the generic approach of using an experience memory to help exploration by adapting a restart distribution. That is, given the capacity to reset the state with those corresponding to the agent’s past observations, we help exploration by promoting faster state-space coverage via restarting the agent from a more diverse set of initial states, as well as allowing it to restart in states associated with significant past experiences. This approach is compatible with both on-policy and off-policy methods. However, a caveat is that altering the distribution of initial states could change the optimal policies when searching within a restricted class of policies. To reduce this unsought learning bias, we evaluate our approach in deep reinforcement learning which benefits from the high representational capacity of deep neural networks. We instantiate three variants of our approach, each inspired by an idea in the context of experience replay. Using these variants, we show that performance gains can be achieved, especially in hard exploration problems.

## 1 Introduction

Experience replay lets off-policy reinforcement learning (RL) methods remember and reuse past experiences [Lin, 1992; Mnih et al., 2015]. This helps circumvent the rapid forgetting of past experiences and, therefore, improves sample-efficiency. Prioritising experience can further boost efficiency by replaying important transitions more frequently, where different criteria may be considered to measure the importance of each transition. For example, the magnitude of a transition’s temporal-difference (TD) error can be used as a proxy for how unexpected the transition is [van Seijen and Sutton, 2013; Schaul et al., 2016]. Transitions can also be rated based on their corresponding episodic return [Oh et al., 2018], a particularly useful criterion in environments with sparse rewards.

On the other hand, on-policy methods cannot benefit from experience replay. As such, they are often sample-inefficient as past transitions are thrown away shortly after they are experienced, regardless of how rare or significant they may be. While replaying past experiences is not compatible with on-policy methods, creating new ones near previously-encountered states is. Given the capacity to reset the state with those corresponding to the agent’s past observations (e.g. in a standard simulator), the latter can be made possible by maintaining a memory of the agent’s previously-encountered states and using it to sample initial states. Effectively, this modifies the perceived distribution of initial states by combining the environment’s *initial-state distribution* with a proposal *restart distribution* over the buffered states, where different criteria could be considered to *prioritise* the latter.

---

\*Equal contribution. Correspondence to: Arash Tavakoli ([a.tavakoli@imperial.ac.uk](mailto:a.tavakoli@imperial.ac.uk)).

We refer to this approach, generically, as *exploring restart distributions*.<sup>2</sup> By drawing inspiration from well-known ideas in the context of experience replay, we instantiate three variants of our approach. Specifically, our *uniform restart* resembles the uniform replay of [Mnih et al., 2015], our *prioritised restart* resembles the prioritised replay of [Schaul et al., 2016], and our *episodic restart* resembles the episodic replay of [Oh et al., 2018]. We combine our variants with a canonical policy-gradient algorithm, Proximal Policy Optimisation (PPO) [Schulman et al., 2017], which, due to its on-policy nature, cannot straightforwardly use experience replay and, as such, is interesting for our study. We test the resulting agents on two dense-reward and two sparse-reward environments, in each case considering a medium-difficulty and a hard exploration problem. We see improvements from our approach in all cases, with the most remarkable gains in the hard exploration problems.

Broadly, we consider *simulator-based training* with *simulator-free execution*, a problem paradigm in which we can utilise the opportunity to adjust certain environment variables during training but not during execution (e.g. [Ciosek and Whiteson, 2017]). Our approach improves this paradigm by utilising the reset capacity in simulated environments during training. We emphasise that we do not utilise this capacity during evaluations (i.e. policies are evaluated with respect to the original performance metric).<sup>3</sup>

## 2 Related work

Kakade and Langford [2002] considered the notion of utilising the reset capacity and showed, under certain conditions, that using a proposal initial-state distribution that is more uniform over the state space than the original one improves learning performance with respect to the original performance metric. More recently, Agarwal et al. [2019] formalised the importance of how a favourable initial-state distribution provides a means to circumvent worst-case exploration issues in the context of policy-gradient methods [Sutton et al., 1999]. Nonetheless, these works do not provide a practical procedure for creating such distributions when the state space is unknown a priori.

To improve model-free learning, Popov et al. [2017] modified the initial-state distribution to be uniform over the states from provided expert demonstrations. Salimans and Chen [2018] reported high performance on Montezuma’s Revenge Atari 2600 game by restarting a standard deep RL agent from a set of designated initial states, manually extracted from a single expert demonstration. These approaches resemble our episodic restart variant with one major difference: in our approach, the agent progressively updates its best buffered episodes in order to sample initial states from them and, as such, it does not rely on expert demonstrations or manually designated initial states.

Ecoffet et al. [2019] proposed a related method, called Go-Explore, that achieved the state-of-the-art on the hardest exploration games in the Atari 2600. Go-Explore’s main principles are to maintain a memory of previously-encountered states, reset the environment to the “promising” ones to explore from, and repeat this process until a complete solution is found. Once found, a policy is trained by imitation learning on the solution trajectory. This work provides strong supporting evidence for the utility of exploration through restarting from previously-encountered states. However, Go-Explore does not use RL to learn a policy that solves the problem. Furthermore, the criterion used to identify promising states is rather domain-specific. Using our approach, one could realise an RL counterpart for Go-Explore by using the same criterion to prioritise initial states.

Florensa et al. [2017] presented a method for adaptive generation of curricula in the form of initial-state distributions that start close to the goal state and gradually move away with the agent’s progress. This method is limited to goal-oriented problems with clear goal states and further assumes a priori knowledge of such states. While our approach is not limited to such environments, a similar behaviour to curriculum generation in this way could emerge with our approach using an appropriate priority measure, whereby a single encounter of a goal state biases the restart distribution towards it.

Restarting from previously-encountered states to sample more transitions reduces the variance of the gradient estimator in policy-gradient methods. The vine procedure of [Schulman et al., 2015]

---

<sup>2</sup>This choice was inspired in part by the theoretical assumption of *exploring starts* [Sutton and Barto, 2018], with which our approach shares a subtle connection.

<sup>3</sup>As such, our work differs from [Rajeswaran et al., 2017] which examines the impact of more diverse initial-state distributions in the context of “robustness”.

utilises the reset capacity in simulated environments for this purpose. This method can be realised as a special case of our approach.

### 3 Background

We consider the RL framework [Szepesvári, 2010; Sutton and Barto, 2018] in which the interaction of an agent and an environment is modeled as a Markov decision process (MDP) [Puterman, 1994] comprising of a state space  $\mathcal{S}$ , an action space  $\mathcal{A}$ , an initial-state distribution  $p_1(s_1) = Pr\{S_1=s_1\}$ , a transition distribution  $p(s'|s, a) = Pr\{S_{t+1}=s'|S_t=s, A_t=a\}$ , and a reward function  $r(s, a, s') = \mathbb{E}[R_t|S_t=s, A_t=a, S_{t+1}=s']$ , for all  $s, s' \in \mathcal{S}, a \in \mathcal{A}, s_1 \in \mathcal{S}_1 \subset \mathcal{S}$ . The decision-making process of an agent is characterised by a policy  $\pi(a|s) = Pr\{A_t=a|S_t=s\}$ . This policy can be approximated by a parameterised function  $\pi(a|s, \theta)$  (e.g. a neural network), where  $\theta \in \mathbb{R}^d$  is the vector of policy parameters and, typically,  $d \ll |\mathcal{S}|$ . The agent uses its policy to interact with the environment to sample a trajectory  $S_1, A_1, R_1, S_2, \dots, S_T, A_T, R_T, S_{T+1}$  (where  $T$  is the trajectory’s horizon which is, in general, a random variable). In this paper, we assume that  $T$  is finite and that terminations may occur due to terminal states in episodic tasks (i.e. *concrete episodes*) or due to an arbitrary condition, such as timeouts, in continuing or episodic tasks (i.e. *partial episodes*). The majority of our discussions are considered under the more generic assumption of learning from partial episodes and, as such, are relevant only to bootstrapping methods [Sutton and Barto, 2018] (e.g. TD methods such as Q-learning, Sarsa, and actor-critic methods). Nevertheless, the main proposition of this paper applies also to Monte-Carlo methods, in which case the episodes are strictly concrete.

We assume access to the capacity to reset the state with those corresponding to the agent’s past observations. We remark that this assumption is weaker than having explicit access to the environment’s model. Furthermore, we do not assume a priori knowledge of the (valid) state space. In fact, such knowledge is rarely accessible in practice, which is why we build a memory of states on-the-fly.

#### 3.1 Impact of the initial-state distribution on the learning objective

In this section, we consider the question “how does modifying the initial-state distribution affect the learning objective and, ultimately, the learned policy with respect to the original performance metric?”. We will consider this question separately for tabular and approximate solution methods.

In tabular methods, the learned values at each state are decoupled from one another (i.e. an update at one state affects no other). Let us now consider the control problem in which the agent’s goal is to maximise its value from the environment’s designated set of initial states. As per the *principle of optimality* [Sutton and Barto, 2018], a policy achieves the optimal value from a state  $s$ , if and only if, for any state  $s'$  reachable from  $s$  it achieves the optimal value. Therefore, by letting the agent also start in states outside the environment’s designated set of initial states, we can better optimise for the designated set by better optimising for the states that are reachable from the designated set.

On the contrary, with approximation, an update at one state affects many others as generally we have far more states than parameters. Therefore, making one state’s estimate more accurate often means making others’ less accurate. Let us now consider a common objective function for approximate prediction of the action values for a given policy:

$$L(\mathbf{w}) \doteq \sum_{s \in \mathcal{S}} \rho(s) \sum_{a \in \mathcal{A}} \pi(a|s) \left( q_\pi(s, a) - \hat{q}_\pi(s, a|\mathbf{w}) \right)^2. \quad (1)$$

This objective function is weighted according to the state distribution  $\rho(s)$ , which depends on the policy  $\pi(a|s)$  and, in episodic tasks, the initial-state distribution  $p_1(s_1)$ . In effect, this promotes the approximation of the action values to be more accurate at states that have a higher visitation density. As such, changing the initial-state distribution modifies the learning objective for approximate prediction. The same rationale holds in approximate control, e.g. when using policy-gradient methods. One caveat to this in the control case is a policy that maximises the modified learning objective within some restricted class of policies may perform poorly with respect to the original performance metric. We can reduce this unsought learning bias by using a distribution of initial states whose support contains and spans beyond that of the environment’s initial-state distribution (see Sec. 4), as well as using a parameterisation that affords the problem’s underlying complexity. While the latter cannot be guaranteed in general, it seems often admissible in deep RL (especially considering the

relative simplicity of many problems of interest with respect to the commonly-used, high-capacity neural networks [Rajeswaran et al., 2017]). Nonetheless, there are many problems where learning any reasonable policy is challenging, not to mention learning an optimal one. In such cases, it may be appropriate to accept the cost of potentially introducing a learning bias in order to facilitate learning.

## 4 Exploring restart distributions

We consider a generic approach in which the agent maintains, what we call, a *restart memory* of its past experiences along with their corresponding (true) states, and uses this restart memory to sample initial states for new episodes. This in turn allows the agent to gradually increase the diversity of the states in which it can restart. Formally, for the environment’s initial-state space  $\mathcal{S}_1$  and the agent’s set of buffered states  $\mathcal{S}_B$ , our approach enables sampling initial states from  $\mathcal{S}_1 \cup \mathcal{S}_B$  (which contains and spans beyond  $\mathcal{S}_1$ ). We achieve this by sampling from both the environment’s initial-state distribution  $p_1$  (with support  $\mathcal{S}_1$ ) and a restart distribution  $p_B$  (with support  $\mathcal{S}_B$ ). This is equivalent to sampling from a new distribution  $\mu_1$ , which mixes  $p_1$  and  $p_B$ , with support  $\mathcal{S}_1 \cup \mathcal{S}_B$ . In this paper, we control the extent of contributions from each of the distributions  $p_1$  and  $p_B$  by maintaining a fixed *ratio* for the number of transitions that stem from *augmented initial states* (i.e. initial states that are sampled from a restart memory) versus the total number of transitions.

### 4.1 Uniform restart

Our uniform restart variant generally follows the same mechanism as the uniform replay of [Mnih et al., 2015]. The differences are that we store states as opposed to observations and that we ultimately care about sampling states as opposed to transitions. In other words, we store recent states (without selection) in a restart memory  $\mathcal{B}$  and sample augmented initial states uniformly (i.e.  $p_B$  is a uniform distribution over the buffered states  $\mathcal{S}_B$ ).

Having the capacity to reset the state naturally implies we can early-terminate episodes. Doing so is compatible with bootstrapping methods which can bootstrap at the end of partial episodes [Pardo et al., 2018]. By convention, we choose to apply a time limit  $T_{\text{aug}}$  to interactions that stem from augmented initial states with  $T_{\text{aug}} \leq T_{\text{env}}$ , where  $T_{\text{env}}$  is the environment’s time limit (if any).

### 4.2 Prioritised restart

Our prioritised restart variant uses a similar mechanism as the *proportional* prioritised replay of [Schaul et al., 2016] but for prioritising states rather than transitions. As such, we use the state-value TD error (as opposed to the state-action form used in [Schaul et al., 2016]):

$$\delta_i \doteq r_i + \gamma v(s'_i) - v(s_i), \tag{2}$$

where  $i$  is the index of state  $s_i$  in the restart memory. We calculate the probability of sampling state  $s_i$  from the restart memory via

$$p_B(s_i) = \frac{p_i^\alpha}{\sum_k p_k^\alpha}, \tag{3}$$

where  $p_i \doteq |\delta_i| + \varepsilon$  is the priority of state  $s_i$  (with  $\varepsilon$  being a small positive constant to ensure non-zero probabilities for all buffered states) and the exponent  $\alpha$  determines how much prioritisation is used. It is noteworthy that prioritising initial states does not introduce a learning bias in the way that prioritisation in the context of replay does. Biasing the replay frequency directly alters the perceived state transition and reward dynamics in stochastic environments. However, this does not apply to our approach as, regardless of how an initial state is sampled, transitions are always sampled from the environment, not replayed from a replay memory. Therefore, we do not need importance sampling corrections as used in [Schaul et al., 2016]. Lastly, similar to our uniform restart variant, we apply a time limit  $T_{\text{aug}}$  to interactions that stem from augmented initial states.

### 4.3 Episodic restart

Episodic return is another criterion for measuring priorities, one that is particularly useful in environments with sparse rewards [Oh et al., 2018]. We build our episodic restart variant to enable using this criterion for prioritisation, leading to a number of differences with respect to our previous variants.

Most fundamentally, states are now buffered at the end of episodes rather than on each transition, and only episodes that obtain a higher undiscounted return in comparison to those already in the restart memory are buffered. In other words, the agent maintains its most rewarding episodes, where the restart-memory size determines the maximum number of episodes in the restart memory at any given time. The episodes are then prioritised according to their corresponding undiscounted returns, with uniform sampling of states from any selected episodes. While it is possible to also prioritise the states in a selected episode (e.g. using TD error), we omit that in this work to simplify our experiments.

Using our episodic restart variant in the manner described above could significantly hurt the agent’s learning performance by biasing its experiences towards specific parts of the state space. For example, consider a multi-goal problem in which the goal is different in every episode and, as such, the goal is part of the state (i.e. assuming *Markov states* [Sutton and Barto, 2018]). Encountering an episode that has an easy goal could result in obtaining a high return in comparison to other episodes. The states from such an episode receive a high priority for being sampled as initial states in new episodes. This leads to frequently experiencing episodes of the same easy goal, thereby quickly filling the restart memory with episodes of a single goal. To address this, we use a nested storage mechanism in which an episode that stems from an environment’s initial state is stored as a *parent episode* (i.e. a standard episode) and an episode that stems from an augmented initial state is stored as a *sub-episode*, where each sub-episode is linked to the parent episode from which its augmented initial state originated. This avoids filling the restart memory with episodes of similar nature as, now, only a limited number of sub-episodes can be stored under each parent episode, with the nature of each parent episode being determined by the environment’s initial-state distribution. We will now discuss what is needed and how to sample an augmented initial state from this restart memory.

To maintain episodic returns comparable in environments with time limits, all sub-episodes need to be terminated after an appropriate number of steps,  $T_{\text{aug}}$ . Assuming (as before) a fixed, environment’s time limit  $T_{\text{env}}$ , sub-episodic time limit  $T_{\text{aug}}$  can be determined via

$$T_{\text{aug}} = T_{\text{env}} - t \quad \text{with } t \leq T_{\text{env}} - 1, \quad (4)$$

where  $t$  is the time step of the augmented initial state.<sup>4</sup> In this way, we can calculate an augmented episodic return for any sub-episodes as the sum of the sub-episode’s rewards and the rewards along the shortest path from the parent episode’s initial state to the sub-episode’s augmented initial state.

To calculate the priority of a *category of episodes* (i.e. a parent episode together with its sub-episodes), we use the maximum undiscounted return across the category’s parent episode and all its sub-episodes. Let us denote this maximum undiscounted return per category with  $\bar{G}_i$ , with  $i$  being the index of the category in the restart memory. We calculate the priority of the  $i$ th category as

$$p_i \doteq \bar{G}_i - \delta_i + \varepsilon, \quad (5)$$

where

$$\delta_i \doteq \min(0, \min_i \bar{G}_i) \quad (6)$$

is used as an offset to enable handling negative returns and  $\varepsilon$  is a small positive constant that ensures non-zero probabilities for all buffered categories. We calculate the probability of sampling the  $i$ th category from the restart memory by using the priorities of Eq. (5) in Eq. (3). We perform this once to sample a category of episodes, and again to sample an episode from the selected category.<sup>5</sup> We then sample an augmented initial state uniformly from the selected episode.

## 5 Experiments

We evaluate our approach on several continuous control environments, simulated using the MuJoCo physics engine [Todorov et al., 2012]. As per the nature of our variants, we present performance evaluations of our uniform and prioritised restart variants in dense-reward environments (Sec. 5.1) and of our episodic restart variant in sparse-reward environments (Sec. 5.2). In each case, we consider a medium-difficulty and a hard exploration problem.

<sup>4</sup>The number of steps in the shortest path from the augmented initial state to the initial state of its corresponding parent episode.

<sup>5</sup>Each sub-episode in a category of episodes is augmented with as many as  $T_{\text{aug}}$  states which link its augmented initial state to the initial state of its corresponding parent episode.

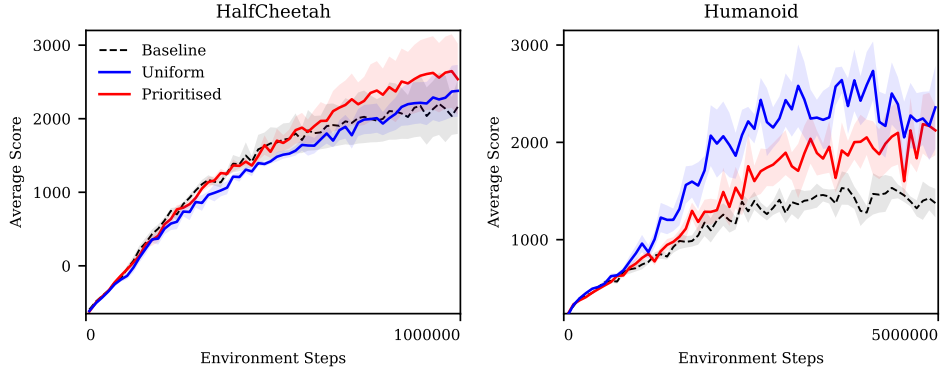


Figure 1: Average test performance curves of our uniform and prioritised restart variants as applied to and against PPO on two dense-reward environments. Shaded areas are standard error.

We focus our experiments on on-policy RL which cannot straightforwardly replay past experiences and, thus, will benefit more significantly from our approach. Specifically, we evaluate our approach using PPO, which is a canonical on-policy method for continuous control. To avoid the significant cost of systematic hyperparameter search, throughout our experiments we fix the ratio hyperparameter of our approach to 0.1 (i.e. 10% of the total interactions stem from augmented initial states) and, generally, use the PPO hyperparameters as originally reported in [Schulman et al., 2017].

## 5.1 Dense-reward environments

To evaluate the performance of our uniform and prioritised restart variants, we consider two dense-reward environments from the OpenAI Gym [Brockman et al., 2016], namely HalfCheetah (medium-difficulty exploration) and Humanoid (hard exploration). By default, these environments apply a time limit of  $T_{\text{env}} = 1000$  to each episode. We apply this default time limit to all interactions that stem from the environment’s initial states. For any interactions that stem from an augmented initial state, we apply the much shorter time limit of  $T_{\text{aug}} = 10$ . For both of our variants, we set the restart-memory size to 20000. For our prioritised restart variant, we set  $\alpha = 0.4$  to induce a mild priority (see Eq. (3)).

### 5.1.1 HalfCheetah

The goal in the HalfCheetah environment is to make a planar biped run as fast as possible. Given the medium-dimensionality of its observation and action spaces, this environment is not very challenging for advanced RL agents. Fig. 1 (left) shows the learning curves for this experiment, created by evaluating each agent periodically during training using the environment’s initial-state distribution. The learning curves show average undiscounted returns (each averaged over 5 seeds). Our results show mild improvements for both of our variants, with a slight advantage for our prioritised one.

### 5.1.2 Humanoid

The goal in the Humanoid environment is to make a three-dimensional biped walk forward as fast as possible, without falling over. While a dense-reward environment, the high-dimensionality of its observation and action spaces make it rather challenging. Fig. 1 (right) shows the learning curves for this experiment, created following the same procedure as in our HalfCheetah experiment. Our results show significant improvements for the agents that use our uniform or prioritised restart variants. However, contrary to our HalfCheetah results, here our uniform restart variant outperforms our prioritised one. We conducted an informal study to examine whether one of these variants would have a general advantage but did not find the results to be conclusive in this regard.

## 5.2 Sparse-reward environments

To evaluate the performance of our episodic restart variant, we consider two sparse-reward environments from the OpenAI Gym, namely FetchReach (medium-difficulty exploration) and Thrower (hard exploration).




Figure 2: Average test success-rate curves of our episodic restart variant as applied to and against PPO on two sparse-reward environments. Shaded areas are standard error.

### 5.2.1 FetchReach

The FetchReach environment was proposed by [Plappert et al. \[2018\]](#) to assess goal-conditional policy learning methods in a problem of practical interest. In each episode, the agent’s task is to control its four joints in order to move the gripper to a goal position, where the goal is randomly sampled at the start of each episode. The agent’s observations consist of the robot’s current state as well as a three-dimensional goal position. This environment has no terminal states but enforces a time limit of  $T_{\text{env}} = 50$ . We set the restart-memory size to 100 parent episodes and 10 sub-episodes. The agent receives a step-wise penalty of -1 whenever its gripper is not at the goal position, and 0 otherwise.

Fig. 2 (left) shows the learning curves for this experiment, created by evaluating each agent periodically during training using the environment’s initial-state distribution. The learning curves show average success rates (each averaged over 10 seeds). Our results show mild improvements for our episodic restart variant. This is consistent with our observation in the HalfCheetah experiment, suggesting that performance improvements can be expected from our approach even in environments which do not pose a significant exploration challenge.

### 5.2.2 Throwing

Our Throwing environment is a variant of the one provided in the OpenAI Gym, modified to feature significantly sparse rewards. Our modifications make this problem particularly challenging as the agent receives a positive reward of 1 only for successfully throwing the ball in the goal region, and 0 otherwise. The agent additionally incurs a small torque penalty. Each episode terminates once the ball impacts the goal or the table, or upon reaching the time limit of  $T_{\text{env}} = 100$ . We set the restart-memory size to 50 parent episodes and 10 sub-episodes. Due to the complexity of this environment, the probability of the ball impacting the goal is very low. Hence, to ensure that each independent run encounters a positive training signal early on in its training process, we only considered runs that experienced an instance of the ball impacting the goal in the first 50000 interaction steps (roughly 10% of our runs achieved this during the specified window). Moreover, we use an entropy coefficient of 0.02 to further encourage exploration.

Fig. 2 (right) shows the learning curves for this experiment, created following the same procedure as in our FetchReach experiment (each averaged over 42 seeds). Our results show a clear advantage for using our episodic restart variant. By inspecting the performances of individual runs, we found that our episodic restart variant enabled the agent to learn consistently across all 42 independent runs (with each run achieving just under 50% success rate per evaluation trial), whereas the baseline agent completely failed to learn any useful policy in 80% of the runs (with the remaining runs achieving just under 50% success rate per evaluation trial, similar to our variant). In other words, using our episodic restart variant led to more robust learning (with respect to random initialisation) by enabling a way for the agent to utilise its extremely rare positive experiences.

## 6 Conclusion

We considered the generic approach of maintaining a restart memory of the agent’s past experiences along with their corresponding (true) states, and using this restart memory to sample initial states for new episodes. This approach utilises the reset capacity in simulated environments during training in order to help with exploration. We instantiated three variants of our approach by drawing inspiration from well-known ideas in the context of experience replay. We tested these variants on two dense-reward and two sparse-reward environments. In each case, we considered a medium-difficulty and a hard exploration problem. We showed improvements from our approach in all cases, with the most remarkable gains in the hard exploration problems.

### Acknowledgements

AT acknowledges financial support from the UK Engineering and Physical Sciences Research Council (EPSRC DTP). VL and PK acknowledge financial support from the Samsung Advanced Institute of Technology (SAIT GRO). All authors acknowledge computational resources from Microsoft (Azure for Research Award).

### References

- Long-Ji Lin. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine Learning*, 8(3):293–321, 1992.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Belle-mare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- Harm van Seijen and Richard S. Sutton. Planning by prioritized sweeping with small backups. In *Proceedings of the International Conference on Machine Learning*, pages 361–369, 2013.
- Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. In *Proceedings of the International Conference on Learning Representations*, 2016.
- Junhyuk Oh, Yijie Guo, Satinder P. Singh, and Honglak Lee. Self-imitation learning. In *Proceedings of the International Conference on Machine Learning*, pages 3878–3887, 2018.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2nd edition, 2018.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv:1707.06347*, 2017.
- Kamil Ciosek and Shimon Whiteson. OFFER: Off-environment reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1819–1825, 2017.
- Aravind Rajeswaran, Kendall Lowrey, Emanuel Todorov, and Sham M. Kakade. Towards generalization and simplicity in continuous control. In *Advances in Neural Information Processing Systems*, pages 6550–6561, 2017.
- Sham M. Kakade and John Langford. Approximately optimal approximate reinforcement learning. In *Proceedings of the International Conference on Machine Learning*, pages 267–274, 2002.
- Alekh Agarwal, Sham M. Kakade, Jason D. Lee, and Gaurav Mahajan. Optimality and approximation with policy gradient methods in Markov decision processes. *arXiv:1908.00261*, 2019.
- Richard S. Sutton, David A. McAllester, Satinder P. Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems*, pages 1057–1063, 1999.



- Ivaylo Popov, Nicolas Heess, Timothy P. Lillicrap, Roland Hafner, Gabriel Barth-Maron, Matej Vecerik, Thomas Lampe, Yuval Tassa, Tom Erez, and Martin Riedmiller. Data-efficient deep reinforcement learning for dexterous manipulation. *arXiv:1704.03073*, 2017.
- Tim Salimans and Richard Chen. Learning Montezuma’s Revenge from a single demonstration. *arXiv:1812.03381*, 2018.
- Adrien Ecoffet, Joost Huizinga, Joel Lehman, Kenneth O. Stanley, and Jeff Clune. Go-Explore: A new approach for hard-exploration problems. *arXiv:1901.10995*, 2019.
- Carlos Florensa, David Held, Markus Wulfmeier, Michael Zhang, and Pieter Abbeel. Reverse curriculum generation for reinforcement learning. In *Proceedings of the Conference on Robot Learning*, pages 482–495, 2017.
- John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *Proceedings of the International conference on machine learning*, pages 1889–1897, 2015.
- Csaba Szepesvári. *Algorithms for Reinforcement Learning*. Morgan & Claypool, 2010.
- Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 1994.
- Fabio Pardo, Arash Tavakoli, Vitaly Levnik, and Petar Kormushev. Time limits in reinforcement learning. In *Proceedings of the International Conference on Machine Learning*, pages 4045–4054, 2018.
- Emanuel Todorov, Tom Erez, and Yuval Tassa. MuJoCo: A physics engine for model-based control. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033, 2012.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI Gym. *arXiv:1606.01540*, 2016.
- Matthias Plappert, Marcin Andrychowicz, Alex Ray, Bob McGrew, Bowen Baker, Glenn Powell, Jonas Schneider, Josh Tobin, Maciek Chociej, Peter Welinder, Vikash Kumar, and Wojciech Zaremba. Multi-goal reinforcement learning: Challenging robotics environments and request for research. *arXiv:1802.09464*, 2018.