
The Primacy Bias in Deep Reinforcement Learning

Evgenii Nikishin^{*1} Max Schwarzer^{*1} Pierluca D’Oro^{*1} Pierre-Luc Bacon¹ Aaron Courville¹

Abstract

This work identifies a common flaw of deep reinforcement learning (RL) algorithms: a tendency to rely on early interactions and ignore useful evidence encountered later. Because of training on progressively growing datasets, deep RL agents incur a risk of overfitting to earlier experiences, negatively affecting the rest of the learning process. Inspired by cognitive science, we refer to this effect as *the primacy bias*. Through a series of experiments, we dissect the algorithmic aspects of deep RL that exacerbate this bias. We then propose a simple yet generally-applicable mechanism that tackles the primacy bias by periodically resetting a part of the agent. We apply this mechanism to algorithms in both discrete (Atari 100k) and continuous action (DeepMind Control Suite) domains, consistently improving their performance.

“Your assumptions are your windows on the world. Scrub them off every once in a while, or the light won’t come in.”

–Isaac Asimov

1. Introduction

Bob is learning a difficult passage on a guitar to rehearse it with his band. After long nights of practice, he is able to play it but with straining finger positions and unclear sound. Alice, another fellow guitarist, shows him a less fatiguing and nicer-sounding way to play the passage which Bob is theoretically able to understand and execute well. Nonetheless, in subsequent rehearsal sessions, Bob’s unconscious mind automatically resorts to the first bad-sounding solution, discouraging him and his bandmates from trying more technically challenging pieces of music. Bob is experiencing an instance of the *primacy bias*, a cognitive bias demonstrated by studies of human learning (Marshall & Werder, 1972; Shteingart et al., 2013).

^{*}Equal contribution ¹Mila, Université de Montréal. Correspondence to: Evgenii Nikishin <evgenii.nikishin@mila.quebec>.

The outcomes of first experiences can have long-lasting effects on subsequent learning and behavior. Thanks to Alice, Bob has already collected new data sufficient for improving his performance on that passage and guitar playing in general; however, because of the priming provided by his long training nights with a bad technique, he is unable to leverage his new experience. The primacy bias generates a vicious circle: since Bob cannot improve his guitar skills in the face of new data, he will not be able to collect more interesting data by playing more challenging guitar passages, crippling his overall learning.

The central finding of this paper is that deep reinforcement learning (RL) algorithms are susceptible to a similar bias. We refer to the primacy bias in deep RL as a tendency to overfit early interactions with the environment preventing the agent from improving its behavior on subsequent experiences. The consequences of this phenomenon compound: an agent with poor performance will collect data of poor quality and the poor data will amplify the difficulty of recovering from an overfitted solution.

Standard components of deep RL algorithms magnify the effect of the primacy bias. For instance, experience replay (Mnih et al., 2015) allows efficient data reuse but exposes the agent to its initial samples more than recent ones. In the interest of sample efficiency, deep RL algorithms often additionally use a high replay ratio (van Hasselt et al., 2019; Fedus et al., 2020) updating the agent more times on the same data. Such design choices can improve the agent’s performance but come with a risk of exacerbating the effects of early interactions.

How can a deep RL algorithm avoid the primacy bias? Coming back to Bob, he could re-establish his learning progression by simply forgetting his bad practices and directly learning from newer experience. Similarly, deep RL agents affected by the primacy bias can forget parts of a solution which was derived by overfitting to early experiences before continuing the learning process.

As a remedy for the primacy bias, we propose a *resetting* mechanism allowing the agent to forget a part of its knowledge. Our strategy is simple and compatible with any deep RL algorithm equipped with a replay buffer: we periodically re-initialize the last layers of an agent’s neural networks, while maintaining the experience within the buffer.

Despite its simplicity, this resetting mechanism consistently improves performance of agents on benchmarks including the discrete-action ALE (Bellemare et al., 2013) and the continuous-action DeepMind Control Suite (Tassa et al., 2020). Our strategy imposes no additional computational costs and requires only two implementation choices: which parts of the neural networks to reset and how often to reset them. We also show that resetting enables training regimes with higher replay ratio and longer n -step targets (Sutton & Barto, 2018), where an agent without resets would be overfitting otherwise.

To summarize the contributions of this paper, we:

1. Provide demonstrations of the existence of the primacy bias in deep RL, a tendency of an agent to harm its future decision making by overfitting to early data and ignoring subsequent interactions;
2. Expose plausible causes of this phenomenon and show how algorithmic aspects in modern deep RL amplify its consequences;
3. Propose a mechanism for alleviating the primacy bias by periodically resetting a part of the agent;
4. Empirically demonstrate both qualitative and quantitative improvements in performance when applying resets to strong baseline algorithms.

2. Preliminaries

We adopt the standard formulation of reinforcement learning (Sutton & Barto, 2018) under the Markov decision process (MDP) formalism where the agent observes a state s from a space \mathcal{S} , chooses an action a from a space \mathcal{A} , and receives a reward r according to a mapping $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. The environment then transitions into a state s' according to a distribution $p : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ and the interaction continues. The initial state s_0 is sampled from a distribution $\rho \in \Delta(\mathcal{S})$. The goal of the agent is to learn a policy $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ that maximizes the expected discounted sum of rewards $\mathbb{E}_\pi [\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t)]$ with $\gamma \in [0, 1)$.

RL methods typically learn an action-value function $Q_\pi(s, a) = \mathbb{E}_\pi [\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) | s_0 = s, a_0 = a]$ through *temporal-difference* (TD) learning (Sutton, 1988) by minimizing the difference between $Q_\pi(s, a)$ and $\mathbb{E}_{p(s'|s,a), \pi(a'|s')} [r(s, a) + \gamma Q_\pi(s', a')]$. Many RL algorithms store past experiences in a *replay buffer* (Mnih et al., 2015) that increases sample efficiency by leveraging a single data point more than once. The number of resampling times of given data is controlled by the *replay ratio* (van Hasselt et al., 2019; D'Oro & Jaśkowski, 2020) which plays a critical role in the algorithm's performance. Set too low, the agent would underuse the data it has and become sample-

inefficient; set too high, the agent would overfit the existing data.

The core idea behind temporal-difference methods is *bootstrapping*, learning from agent's own value estimates without the need to wait until the end of the interaction. TD learning can be generalized by using n -step targets $\mathbb{E}_\pi [r(s_t, a_t) + \gamma r(s_{t+1}, a_{t+1}) + \dots + \gamma^n Q_\pi(s_{t+n}, a_{t+n})]$. Here, n controls a trade-off between the (statistical) bias of Q_π estimates and the variance of the sum of future rewards.

In the rest of the paper, we consider deep RL algorithms where Q_π and π (when needed) are modelled by neural network function approximators.

3. The Primacy Bias

The main goal of this paper is to understand how the learning process of deep reinforcement learning agents can be disproportionately impacted by initial phases of training due to an effect called the primacy bias.

The Primacy Bias in Deep RL: *a tendency to overfit early experiences that damages the rest of the learning process.*

This definition is wide-ranging: the primacy bias has multiple roots and leads to multiple negative effects for the training of an RL agent, but they are all connected to improper learning from early data.

The rest of this section presents two experiments, with the goal of demonstrating the existence and the dynamics of the phenomenon in isolation. First, we show that excessive training of an agent on the very first interactions can fatally damage the rest of the learning process. The second experiment demonstrates that data collected by an agent impacted by the primacy bias is adequate for learning, although the agent cannot leverage it due to its accumulated overfitting.

3.1. Heavy Priming Causes Unrecoverable Overfitting

The degree of reliance of an agent on early data is a crucial factor in determining how much any primacy effect is going to affect the learning process. At the same time, in the interest of sample efficiency, it is vital to leverage initial experiences at their full potential to expedite the training. This trade-off is particularly evident for algorithms with a replay buffer, which can be used to update the agent several times before interacting further with the environment.

To uncover in an explicit way the effect of the primacy bias in RL, we probe excessive reliance to early data to its extreme: could *overfitting on a single batch of early data be enough to entirely disrupt an agent's learning process?*

To investigate this question, we train Soft Actor-Critic (Haarnoja et al., 2018) on the quadruped-run environment from DeepMind Control suite (DMC) (Tassa




Figure 1. Undiscounted returns on quadruped-run for SAC with and without *heavy priming* on the first 100 transitions. An agent extremely affected by the primacy bias is unable to learn even after collecting hundreds of thousands of new transitions. Mean and std are estimated over 10 runs.

et al., 2020). We use default hyperparameters, which imply a single update for both policy and value function per step in the environment. Then, we train an identical version of the algorithm in an experimental condition that we refer to as *heavy priming*: after collecting 100 data points, we update the agent 10^5 times using the resulting replay buffer, before resuming standard training. Figure 1 shows that even after collecting and training on almost one million new transitions, the agent with heavy priming is unable to recover from the initial overfitting.

This experiment conveys a simple message: overfitting to early experiences might inexorably damage the rest of the learning process. Even if no practical implementation would use such a large number of training steps on such a limited dataset, we will see in Section 5 that even a relatively small number of updates per step can cause similar issues. The finding suggests that the primacy bias has compounding effects: an overfitted agent gathers worse data that itself leads to less efficient learning that further damages the ability to learn and so on.

3.2. Experiences of Primed Agents are Sufficient

Once the agent is heavily impacted by the primacy bias, it might struggle to reach satisfying performance. But is the data collected by an overfitted agent unusable for learning? To answer this question, we train a SAC agent with 9 updates per step in the MDP: due to the primacy bias, this agent performs poorly. Then, we initialize the same agent from scratch but use the data collected by the previous SAC agent as its initial replay buffer. Figure 2 demonstrates that returns collected by this agent improve rapidly approaching the




Figure 2. Undiscounted returns on quadruped-run for SAC trained with 9 updates per step. SAC failing is a standard agent; SAC with failing agent buffer is an agent initialized with the replay buffer of the first agent, which allows it to learn quickly. Mean and std are estimated over 10 runs.

optimal task performance.

This experiment articulates that *the primacy bias is not a failure to collect proper data per se, but rather a failure to learn from it*. The data stored in the replay buffer is in principle enough to have better performance but the overfitted agent lacks the ability to distill it into a better policy. In contrast, the randomly initialized neural networks are not affected by the primacy bias and thus capable of fully leveraging the collected experience.

4. Have You Tried Resetting It?


The previous section provided controlled experiments demonstrating the primacy bias phenomenon and outlined its consequences. We now present a simple technique that mitigates the effect of this bias on an agent’s training. The solution, which we call *resetting* in the rest of the paper, is summarized as follows:

Given an agent’s neural network, periodically re-initialize the parameters of its last few layers while preserving the replay buffer.

The next section analyzes both quantitatively and qualitatively the performance improvements provided by resetting as a mean to address the overfitting to early data.

5. Experiments

The goals of our experiments are mostly twofold. First, we demonstrate across different algorithms and domains the



Method	IQM	Median	Mean
SPR + resets	0.478 (0.46, 0.51)	0.512 (0.42, 0.57)	0.911 (0.84, 1.00)
SPR	0.380 (0.36, 0.39)	0.433 (0.38, 0.48)	0.578 (0.56, 0.60)
DrQ(ϵ)	0.280 (0.27, 0.29)	0.304 (0.28, 0.33)	0.465 (0.46, 0.48)
DER	0.183 (0.18, 0.19)	0.191 (0.18, 0.21)	0.351 (0.34, 0.36)
CURL	0.113 (0.11, 0.12)	0.102 (0.09, 0.12)	0.261 (0.25, 0.27)

Figure 3 & Table 1. Point estimates and 95% bootstrap confidence intervals for the performance of SPR with resets and prior methods on Atari 100k. Results for SPR and SPR + resets are over 20 seeds per game; others are taken from Agarwal et al. (2021) and use 100 seeds. The additional baselines contextualize the numerical impact of resets. Per-environment results are available in Appendix C.

performance gains of using resets as a remedy for the primacy bias; then, we analyze the learning dynamics induced by resetting, including its effects on TD learning and interaction with critical design choices of RL algorithms such as the replay ratio and n -step TD targets. Finally, we provide an ablation study of the proposed resetting strategy.

5.1. Setup

We focus on two domains: discrete control, represented by the 26-task Atari 100k benchmark (Kaiser et al., 2019), and continuous control, represented by the DeepMind Control Suite (Tassa et al., 2020). We apply resets to three baseline algorithms: SPR (Schwarzer et al., 2020) for Atari, and SAC (Haarnoja et al., 2018) and DrQ (Kostrikov et al., 2021) for continuous control from dense states and raw pixels respectively. Appendix A provides all environment names and the number of training steps in each domain.

Since both the architectures and the number of training iterations vary across methods, the reset strategy needs slight customization. For SPR, we reset only the final linear layer of the 5-layer Q-network over the course of training spaced 2×10^4 steps apart; for SAC, we reset agent’s networks entirely every 2×10^5 steps since the networks have only 3 layers; for DrQ, we reset the last 3 out of 7 layers of the policy and value networks 10 times over the course of training¹. SAC and DrQ re-initialize target networks and both Q-networks (due to the use of double Q-learning (Van Hasselt et al., 2016)); SPR does not have these extra networks. We also reset the corresponding optimizer statistics (Kingma & Ba, 2015). Appendix B, however, shows the relative robustness to these design choices.

The replay buffer is preserved between resets; SPR and SAC store in the buffer all prior interactions, while DrQ includes only the most recent 100k transitions due to memory limitations of storing image observations. SAC and DrQ sample

¹In fact, the reset periodicity in a few environments is higher due to action repeats (default is 2), a practice following (Hafner et al., 2019) used in the codebase we build upon, but using the action repeat of 2 for all environments delivers the same results.

Method	IQM	Median	Mean
SAC + resets	656 (549, 753)	617 (538, 681)	607 (547, 667)
SAC	501 (389, 609)	475 (407, 563)	484 (420, 548)
DrQ + resets	762 (704, 815)	680 (625, 731)	677 (632, 720)
DrQ	569 (475, 662)	521 (470, 600)	535 (481, 589)

Table 2. Point estimates and 95% bootstrap confidence intervals for the performance of SAC and DrQ with and without resets on DMC tasks. Results are computed over 10 seeds per task. Per-environment learning curves are available in Appendix C.

transitions uniformly from the buffer, while SPR uses prioritized experience replay (Schaul et al., 2016). The difference between buffer configurations suggests that effects of resets hold for varying buffer sizes and sampling schemes.

After resetting, we do not perform any form of pre-training for the newly initialized parameters (Igl et al., 2021) and return directly to standard training, including keeping intact the cycle between environment interactions and agent updates. To provide rigorous evaluation of all algorithms, we follow the guidelines of Agarwal et al. (2021) with the focus on the interquartile mean (IQM) of the performance across tasks.

5.2. Resets Consistently Improve Performance

Tables 1 and 2 report the aggregated results for the three algorithms. In both tables, we report the best results over different values of replay ratio and n for methods with and without resets. The empirical evidence suggests that resets mitigate the primacy bias and provide significant benefits across a wide range of tasks (discrete or continuous action spaces), input types (raw images or dense features), replay buffer configurations (matching or shorter than total number of steps, prioritized replay or random sampling), and depth of the employed neural networks (deep convolutional networks or 3-layer fully-connected networks). Remarkably, the magnitude of improvement provided by resets for SPR is comparable to advancements of prior algorithms, while not requiring additional computation costs.

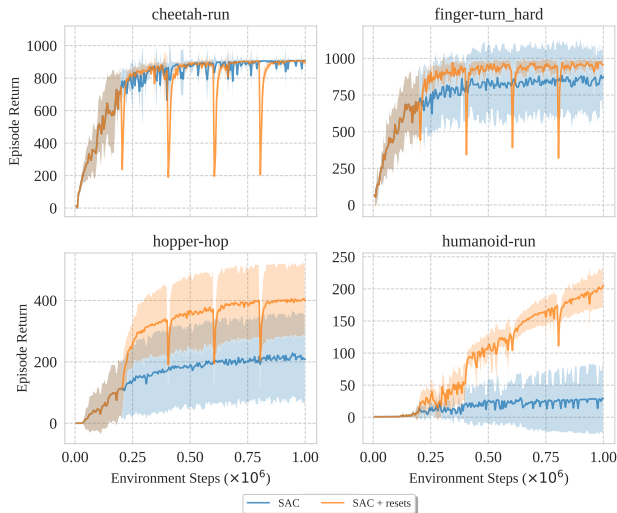


Figure 4. Four examples showing diverse effects of resets for SAC (32 updates per step, resetting every 2×10^5 steps) on DMC tasks. After each reset, performance recovers quickly due to keeping the replay buffer. In `cheetah-run`, the baseline agent consistently succeeds at the task and resets provide no major benefit. In all other tasks, resets increase performance and often reduce variance. Mean and std are estimated over 10 runs.

5.3. Learning Dynamics of Agents with Resets

At first glance, resetting may appear as a drastic (if not wasteful) measure as the agent must learn the parameters of the randomly initialized layers from scratch each time. We show in this section how, against all odds, this strategy still leads to improved performance and fast learning in a wide range of situations.

Figure 4 gives four representative examples of the learning trajectories induced by resets. By default, SAC uses policy and value function architectures which typically contain three layers; in this context, we found that resetting *the whole network* is an effective strategy. For environments in which the primacy bias does not appear to be an issue, such as `cheetah-run`, resetting causes some spikes of reduced performance in the learning curves, but the agent is able to return to the previous state in just a few thousands of steps. Instead, when dealing with environments where the algorithm is susceptible to the primacy bias, such as `hopper-hop` and `humanoid-run`, resetting not only brings performance back to the previous level but also allows the agent to surpass it.

But why is an RL agent able to recover so fast after each reset? A decisive factor for the effectiveness of resetting resides in preserving the replay buffer across iterations. Indeed, periodically emptying the replay buffer is highly detrimental for performance, as we show in Appendix B. We conjecture that a model-based perspective can offer an ex-

planation: since the replay buffer can be seen as a non-parametric model of the world (Vanseijen & Sutton, 2015; van Hasselt et al., 2019), after a reset, the agent forgets the behaviors learned in the past while preserving its model in the buffer as the core of its knowledge. On the neural network training side, Zhang et al. (2019) observe that learning mostly amounts to recovering the right representations – that is, with the preserved buffer and representations, learning an actuator might be relatively straightforward.

Resets affect learning in a generally positive way, by triggering a virtuous circle. After resetting, the agent is free from the negative priming provided by its past training iterations: it can better leverage the data collected so far, thus improving its performance and unlocking the possibility to generate higher quality data for its future updates.

If the primacy bias is a special form of overfitting, resets can be seen as a tailor-made form of regularization. Table 5 in Appendix B shows that the particular nature of resets allows the agent to overcome the primacy bias even when other forms of regularization such as L2 and dropout would not.

On a practical side, resets are an easy-to-use strategy for addressing the primacy bias. Their use requires making only two choices: the periodicity of the resets and the number of layers of the neural networks to be reinitialized. Moreover, the infrequent re-initialization of a neural network comes with no additional computational cost.

5.4. Elements Behind the Success of Resets

The large improvement in performance provided by resets across algorithms and environments naturally raises a question about the conditions under which they are maximally useful. To find an answer, we focus the discussion on the interaction of resets with crucial algorithmic aspects and hyperparameters impacting the risk of overfitting.

Replay ratio The initial experiments in Section 3 suggest that the degree of reliance on early data is a critical determinant of the strength of the primacy bias. As a consequence, we observe that the impact of resets depends heavily on the *replay ratio*, the number of gradient updates per each environment step. Figure 5 shows that the higher the replay ratio is, the larger the effects from resets: they improve SPR’s performance by over 40% at four updates/step and allow SAC to achieve its highest performance at the high replay ratio of 32, where adding resets increases performance by over 100%. Even when pushing SAC to the extreme replay ratios of 128 and 256, where learning is barely possible in most environments, resets allow the agent to obtain reasonable performance (see Table 6 in Appendix C). Resets thus allow less careful tuning of this parameter and improve sample efficiency by performing more updates per each data point without being severely affected by the primacy bias.




Figure 5. Performance of SAC (left) and SPR (right) and with and without resets for different replay ratios and a fixed default n . The right-hand plots visualize the percent improvement gained by adding resets. Agents with higher replay ratio are more prone to the primacy bias and hence benefit more from mitigating it.

n -step targets The parameter n in TD learning controls a bias-variance trade-off, with larger values of n decreasing the bias in value estimates but increasing the variance (and vice versa). We hypothesize that an agent learning from higher variance targets would be more prone to overfitting to the initial data and the effects of resets would increase with increasing n . Figure 6 confirms the intuition and demonstrates up to 40% improvement for SPR for $n = 20$ and opposed to no improvement for $n = 3$. Likewise, SAC attains 50–60% improvement for increased values of n compared to 40% for the default $n = 1$.

The results with varying replay ratios and n -step targets suggest that resets reshape the hyperparameter landscape creating a new optimum with higher performance.

TD failure modes Temporal-difference learning, when employed jointly with function approximation and off-policy training, is known to be potentially unstable (Sutton & Barto, 2018). In sparse-reward environments, the critic network might converge to a degenerate solution because of bootstrapping mostly on its own outputs (Kumar et al., 2020); having the non-zero reward data might not sufficient to escape a collapse. For example, Figure 7 (left) demonstrates the behavior of DrQ in `cartpole-swingup-sparse`, where a collapsed agent makes no learning progress. However, after a manual examination of a replay buffer, we found that the agent was reaching goal states in roughly 2% of trajectories. This ob-

ervation provides evidence for an explanation that *mitigating the primacy bias addresses more the issues of optimization than exploration*. Likewise, if divergence in temporal difference learning occurs, it is essentially unrecoverable by standard optimization, with predicted values failing to decay to normal magnitudes after hundreds of thousands of steps. Figure 7 shows that adding resets solves this problem by giving the agent a second chance to find a stable solution. Even though there exist works studying in detail the pathological behaviors of TD learning (Bengio et al., 2020) and this is not the main scope of our work, resets naturally address the outlined failures.

What and how to reset Our particular choice of the resetting strategy calls for a number of ablations. This paragraph provides only the conclusions while Appendix B presents supporting figures. The number of layers to reset is a domain-dependent choice. For the SAC algorithm learning from dense state features, it is possible to reset the agent entirely. Resetting in SPR attains the best performance for the last layer only (out of 5 total layers), while for DrQ resetting the last layer is slightly worse than resetting last 3 out of 7 layers. We conjecture that the difference lies in the degree of representation learning required for each domain: a significant chunk of knowledge in Atari is contained in the agent’s representations; it might be notably easier to learn features in DeepMind Control, especially when dealing with dense states. In DrQ, when resetting both actor and critic,




Figure 6. Performance of SAC (left) and SPR (right) with and without resets for different n -step target lengths and a fixed replay ratio (9 for SAC, default 2 for SPR). The right-hand plots visualize the percent improvement gained by adding resets. As the target variance increases with n , the agent becomes more susceptible to the primacy bias and benefits more from mitigating it.




Figure 7. Examples of TD failure modes and how resets address them. **Left:** A run with TD collapse in a sparse-reward task `cartpole-swingup-sparse`. Even in the presence of non-zero rewards in the buffer, the agent without resets cannot learn a non-trivial critic. **Right:** A run with TD divergence in `walker-stand`. Even with double Q-learning, the critic might severely overestimate the action values. On both plots, DrQ without resets achieves near-zero returns, while DrQ + resets learns a near-optimal policy. The examples are not cherry-picked, such patterns of behavior occur frequently.

resetting critic proved to be slightly more important than the actor; likely because the DrQ encoder learns from critic loss only, a practice proposed in (Yarats et al., 2021).

Another seemingly important choice is whether to reset the state of the optimizer. We find that resetting the optimizer has almost no impact on training because the moment estimates are updated quickly. Regarding the resets frequency, the optimal choice should depend on how fast an algorithm can recover and how much it is affected by the primacy bias; we found that sometimes even a single reset improves the performance of baseline agents. We briefly experimented with resetting a random subnetwork and observed that the performance was either comparable or worse than with resetting the last layers. Lastly, when sampling new weights after a reset, it is natural to use a new random seed; we observed that even with the same seed resets alleviate the primacy bias supporting the conclusions of Bjorck et al. (2022) that pathologies in deep RL algorithms are not due to problems with the initialization. Overall, while we see certain differences when varying the discussed design choices, resetting showed itself to be robust the choice of hyperparameters.

5.5. Summary

In short, the experimental results suggest that resets improve expected returns across a diverse set of environments and algorithms. They act as a form of regularization thus preventing overfitting to early data, unlock new hyperparameter configurations with possibly higher performance and sample efficiency, and address the optimization challenges arising in deep RL. While with a more thorough hyperparameter search and additional modifications it is possible to even further improve the performance upon the baselines, it is exciting to see that the proposed simple resetting scheme gives benefits comparable to previous algorithmic advancements.

6. Related Work

The primacy bias in deep RL is intimately related to memorization, optimization in RL, and cognitive science. Various aspects of our work existed in the literature.

Overfitting in RL Generalization and overfitting have many faces in deep reinforcement learning. Generalization of values to similar states is a setting where the most classical form of overfitting can arise when using function approximation (Sutton & Barto, 2018). Kumar et al. (2020) and Lyle et al. (2022) show that an approximator for value function gradually loses its expressivity due to bootstrapping in TD learning; we conjecture that this amplifies the effect of first data points. Dabney et al. (2021) propose to treat holistically the sequence of value prediction tasks, arguing that if an agent focuses too much on a single prediction problem might overspecialize learned representations. Laroche & Des Combes (2022) show that policy gradient updates can be slow in unlearning previous behaviors. Song et al. (2019) study observational overfitting by examining saliency maps in visual domains and argue that agents might pick up spurious correlations for decision making. The pioneering work (Farahmand et al., 2008) and a more recent one (Liu et al., 2021) argue in favor of using regularization in RL. Finally, overfitting can happen in settings with learning from offline datasets (Fujimoto et al., 2019) and with multiple tasks (Teh et al., 2017). Surveys by Kirk et al. (2021) and Zhang et al. (2018) give a taxonomy of generalization aspects in RL.

Techniques similar to our resets also existed before. Anderson et al. (1993) propose to reset individual neurons of a Q-network based on a variance criterion and observes faster convergence. Forms of non-uniform sampling including re-weighting recent samples (Wang et al., 2020) and prioritized experience replay (PER) (Schaul et al., 2016) can be seen as a way to mitigate the primacy bias. We observe that SPR, built on top of Rainbow (Hessel et al., 2018) and already using PER, still benefits from resets. Igl et al. (2021) provide demonstrations in the supervised case that learning from a pretrained network can be worse than learning from scratch for non-stationary datasets and propose a method ITER that fully resets the agent’s network in an on-policy buffer-free setting with distillation from the previous generation as a knowledge transfer mechanism. The difference contrasts the approach with our resetting scheme that uses a replay buffer as a basis for knowledge transfer after a reset.

Our work sheds light on another special form of overfitting in deep RL and proposes a simple solution for mitigating it.

Forgetting mechanisms In contrast to the well-known phenomenon of catastrophic forgetting (French, 1999), several works have noted the opposite effect of catastrophic memorization (Robins, 1993; Sharkey & Sharkey, 1995)

similar to the primacy bias. Achille et al. (2018) observe the existence of critical phases in learning that have a determining effect on the resulting network. Erhan et al. (2010) notice higher sensitivity of the trained network with respect to first datapoints. While the effect of the early examples in supervised problems might be present, the consequences of overfitting to initial experiences in deep RL would be much more drastic because the agent itself collects the data to learn from. Dohare et al. (2021) adjust stochastic gradient descent for the continual learning setting; we highlight that in continual learning the agent does not have influence over the stream of data, while the primacy bias in deep RL arises because data collection is in the training loop.

The idea of resetting subnetworks recently received more attention in supervised learning. Taha et al. (2021) studies this process from an evolutionary perspective and shows increased performance in computer vision tasks. Zhou et al. (2022) show that some degree of forgetting might improve generalization and draw a connection to the emergence of compositional representations (Ren et al., 2019). Zhang et al. (2019) demonstrate that resetting different layers affect differently the performance of a network. Alabdulmohsin et al. (2021) demonstrate that resets increase margins of training examples and induce convergence to flatter minima.

These works complement the evidence about the existence of the primacy bias in deep RL and add to our analysis of the regularization effect of resets.

Cognitive science The primacy bias (also known as the *primacy effect*) is a well-studied cognitive bias in human learning (Marshall & Werder, 1972). Given a sequence of facts, humans often form generalizations based on the first ones and pay less attention to the later ones. Asch (1961) shows that this tendency can foster a creation of harmful prejudices by examining the difference in responses after presenting the same data but in different order. Shteingart et al. (2013) argue that outcomes of first experiences affect future decision making and have a substantial and lasting effect on subsequent behavior. Yalnizyan-Carson & Richards (2021) use RL as a framework to test a hypothesis that some degree of forgetting in natural brains is beneficial for decision making. Resets can be linked to cultural transmission between generations (Kirby, 2001), where an agent before a reset transmits its knowledge to an agent after the reset through a buffer. Lastly, studies of a critical period (Johnson & Newport, 1989) show that if a child fails to develop a skill during a particular stage of its development, it might much more difficult to acquire the skill later, drawing the connection to proper learning from early experiences.

Even though humans and RL systems learn under different conditions, our paper provides evidence that artificial agents exhibit the primacy bias noted in humans.

7. Future Work and Limitations

This paper focuses on empirical investigation of the primacy bias phenomenon. An exciting avenue for future work is developing theoretical understanding of risks associated to overfitting to first experiences. Likewise, deriving guarantees for learning with resets similarly to the results of Li et al. (2021) in games and Besson & Kaufmann (2018) for bandits would make the technique more theoretically sound.

Our version of resets is appealing because of its simplicity. However, the reset periodicity is a hyperparameter that an RL practitioner needs to choose. A version of the technique based on the feedback from the RL system or even meta-learning the resetting strategy can potentially improve the performance even further.

Finally, we note that brief collapses in performance induced by resetting may be undesirable from a regret minimization perspective. Potential remedies include having a period of offline post-training after each reset or sampling actions from an interpolation between pre- and post-reset agents for some period of time after each reset.

8. Conclusion

This paper identifies the primacy bias in deep RL, a damaging tendency of artificial agents to overfit early experiences. We demonstrate dangers associated with this form of overfitting and propose a simple solution based on resetting a part of the agent. The experimental evidence across domains and algorithms suggests that resetting is an effective and generally applicable technique in RL.

The general trend in RL research for many years was to first establish a sound algorithm for the tabular case and then use a neural network for representing parts of the agent. The last step was not rarely seen as just a technical detail. The primacy bias, however, is a phenomenon specific to RL with function approximation. The findings of our paper point at the importance of studying the profound interaction of reinforcement and deep learning. Similarly to techniques like Batch Normalization (Ioffe & Szegedy, 2015) that revolutionized training of supervised models, future progress might come not only from advancements in core RL but rather by approaching the problem in conjunction. It is striking that something as simple as periodic resetting improves performance so drastically, suggesting that there is still much to understand about the behavior of deep RL.

Overall, this work sheds light on the learning processes of deep RL agents, unlocks training regimes that were unavailable without resets, and opens possibilities for further studies improving both understanding and performance of deep reinforcement learning algorithms.

Acknowledgements

The authors thank Rishabh Agarwal, David Yu-Tung Hui, Ilya Kostrikov, Ankit Vani, Zhixuan Lin, Tristan Deleu, Wes Chung, Mandana Samiei, Hattie Zhou, Marc G. Bellemare for insightful discussions and useful suggestions on the early draft, Compute Canada for computational resources, and Sara Hooker for Isaac Asimov’s quote. This work was partially supported by Facebook CIFAR AI Chair, Samsung, Hitachi, IVADO, and Gruppo Ermenegildo Zegna.

We acknowledge the Python community (Van Rossum & Drake Jr, 1995; Oliphant, 2007) for developing the core set of tools that enabled this work, including JAX (Bradbury et al., 2018; Babuschkin et al., 2020), Jupyter (Kluyver et al., 2016), Matplotlib (Hunter, 2007), numpy (Oliphant, 2006; Van Der Walt et al., 2011), pandas (McKinney, 2012), and SciPy (Jones et al., 2014).

References

- Achille, A., Rovere, M., and Soatto, S. Critical learning periods in deep networks. In *International Conference on Learning Representations*, 2018.
- Agarwal, R., Schwarz, M., Castro, P. S., Courville, A., and Bellemare, M. G. Deep reinforcement learning at the edge of the statistical precipice. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.
- Alabdulmohsin, I., Maennel, H., and Keysers, D. The impact of reinitialization on generalization in convolutional neural networks. *arXiv preprint arXiv:2109.00267*, 2021.
- Anderson, C. W. et al. Q-learning with hidden-unit restarting. *Advances in Neural Information Processing Systems*, pp. 81–81, 1993.
- Asch, S. E. *Forming impressions of personality*. University of California Press, 1961.
- Babuschkin, I., Baumli, K., Bell, A., Bhupatiraju, S., Bruce, J., Buchlovsky, P., Budden, D., Cai, T., Clark, A., Danihelka, I., Fantacci, C., Godwin, J., Jones, C., Hennigan, T., Hessel, M., Kapturowski, S., Keck, T., Kemaev, I., King, M., Martens, L., Mikulik, V., Norman, T., Quan, J., Papamakarios, G., Ring, R., Ruiz, F., Sanchez, A., Schneider, R., Sezener, E., Spencer, S., Srinivasan, S., Stokowiec, W., and Viola, F. The DeepMind JAX Ecosystem, 2020. URL <http://github.com/deepmind>.
- 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.
- Bengio, E., Pineau, J., and Precup, D. Interference and generalization in temporal difference learning. In *International Conference on Machine Learning*, pp. 767–777. PMLR, 2020.
- Besson, L. and Kaufmann, E. What doubling tricks can and can’t do for multi-armed bandits. *arXiv preprint arXiv:1803.06971*, 2018.
- Bjorck, J., Gomes, C. P., and Weinberger, K. Q. Is high variance unavoidable in rl? a case study in continuous control. *arXiv preprint arXiv:2110.11222*, 2021.
- Bjorck, J., Gomes, C. P., and Weinberger, K. Q. Is high variance unavoidable in RL? a case study in continuous control. In *International Conference on Learning Representations*, 2022.
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- Dabney, W., Barreto, A., Rowland, M., Dadashi, R., Quan, J., Bellemare, M. G., and Silver, D. The value-improvement path: Towards better representations for reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 7160–7168, 2021.
- Dohare, S., Mahmood, A. R., and Sutton, R. S. Continual backprop: Stochastic gradient descent with persistent randomness. *arXiv preprint arXiv:2108.06325*, 2021.
- D’Oro, P. and Jaśkowski, W. How to learn a useful critic? model-based action-gradient-estimator policy optimization. *Advances in Neural Information Processing Systems*, 33, 2020.
- Erhan, D., Courville, A., Bengio, Y., and Vincent, P. Why does unsupervised pre-training help deep learning? In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pp. 201–208. JMLR Workshop and Conference Proceedings, 2010.
- Farahmand, A., Ghavamzadeh, M., Mannor, S., and Szepesvári, C. Regularized policy iteration. *Advances in Neural Information Processing Systems*, 21, 2008.
- Fedus, W., Ramachandran, P., Agarwal, R., Bengio, Y., Laroche, H., Rowland, M., and Dabney, W. Revisiting fundamentals of experience replay. In *International Conference on Machine Learning*, pp. 3061–3071. PMLR, 2020.
- French, R. M. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4):128–135, 1999.

- Fujimoto, S., Meger, D., and Precup, D. Off-policy deep reinforcement learning without exploration. In *International Conference on Machine Learning*, pp. 2052–2062. PMLR, 2019.
- Haarnoja, T., Zhou, A., Hartikainen, K., Tucker, G., Ha, S., Tan, J., Kumar, V., Zhu, H., Gupta, A., Abbeel, P., et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.
- Hafner, D., Lillicrap, T., Fischer, I., Villegas, R., Ha, D., Lee, H., and Davidson, J. Learning latent dynamics for planning from pixels. In *International Conference on Machine Learning*, pp. 2555–2565. PMLR, 2019.
- Hessel, M., Modayil, J., Van Hasselt, H., Schaul, T., Ostrovski, G., Dabney, W., Horgan, D., Piot, B., Azar, M., and Silver, D. Rainbow: Combining improvements in deep reinforcement learning. In *Thirty-second AAAI conference on artificial intelligence*, 2018.
- Hunter, J. D. Matplotlib: A 2d graphics environment. *IEEE Annals of the History of Computing*, 9(03):90–95, 2007.
- Igl, M., Farquhar, G., Luketina, J., Boehmer, W., and Whiteson, S. Transient non-stationarity and generalisation in deep reinforcement learning. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=Qun8fv4qSby>.
- Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pp. 448–456. PMLR, 2015.
- Johnson, J. S. and Newport, E. L. Critical period effects in second language learning: The influence of maturational state on the acquisition of english as a second language. *Cognitive psychology*, 21(1):60–99, 1989.
- Jones, E., Oliphant, T., and Peterson, P. *SciPy: Open source scientific tools for Python*. 2014.
- Kaiser, Ł., Babaeizadeh, M., Miłoś, P., Osiński, B., Campbell, R. H., Czechowski, K., Erhan, D., Finn, C., Koza-kowski, P., Levine, S., et al. Model based reinforcement learning for atari. In *ICLR*, 2019.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *ICLR (Poster)*, 2015.
- Kirby, S. Spontaneous evolution of linguistic structure-an iterated learning model of the emergence of regularity and irregularity. *IEEE Transactions on Evolutionary Computation*, 5(2):102–110, 2001.
- Kirk, R., Zhang, A., Grefenstette, E., and Rocktäschel, T. A survey of generalisation in deep reinforcement learning. *arXiv preprint arXiv:2111.09794*, 2021.
- Kluyver, T., Ragan-Kelley, B., Pérez, F., Granger, B. E., Bussonnier, M., Frederic, J., Kelley, K., Hamrick, J. B., Grout, J., Corlay, S., et al. *Jupyter Notebooks-a publishing format for reproducible computational workflows.*, volume 2016. 2016.
- Kostrikov, I. JAXRL: Implementations of Reinforcement Learning algorithms in JAX, 10 2021. URL <https://github.com/ikostrikov/jaxrl>.
- Kostrikov, I., Yarats, D., and Fergus, R. Image augmentation is all you need: Regularizing deep reinforcement learning from pixels. In *ICLR*, 2021.
- Kumar, A., Agarwal, R., Ghosh, D., and Levine, S. Implicit under-parameterization inhibits data-efficient deep reinforcement learning. In *International Conference on Learning Representations*, 2020.
- Laroche, R. and Des Combes, R. T. Beyond the policy gradient theorem for efficient policy updates in actor-critic algorithms. In *International Conference on Artificial Intelligence and Statistics*, pp. 5658–5688. PMLR, 2022.
- Li, C. J., Yu, Y., Loizou, N., Gidel, G., Ma, Y., Roux, N. L., and Jordan, M. I. On the convergence of stochastic extragradient for bilinear games with restarted iteration averaging. *arXiv preprint arXiv:2107.00464*, 2021.
- Liu, Z., Li, X., Kang, B., and Darrell, T. Regularization matters in policy optimization - an empirical study on continuous control. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=yrlmzrH3IC>.
- Lyle, C., Rowland, M., and Dabney, W. Understanding and preventing capacity loss in reinforcement learning. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=ZkC8wKoLbQ7>.
- Marshall, P. H. and Werder, P. R. The effects of the elimination of rehearsal on primacy and recency. *Journal of Verbal Learning and Verbal Behavior*, 11(5):649–653, 1972.
- McKinney, W. *Python for data analysis: Data wrangling with Pandas, NumPy, and IPython*. ” O’Reilly Media, Inc.”, 2012.
- 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): 529–533, 2015.
- Oliphant, T. E. *A guide to NumPy*, volume 1. Trelgol Publishing USA, 2006.

- Oliphant, T. E. Python for scientific computing. *Computing in Science & Engineering*, 9(3):10–20, 2007.
- Ren, Y., Guo, S., Labeau, M., Cohen, S. B., and Kirby, S. Compositional languages emerge in a neural iterated learning model. In *International Conference on Learning Representations*, 2019.
- Robins, A. Catastrophic forgetting in neural networks: the role of rehearsal mechanisms. In *Proceedings 1993 The First New Zealand International Two-Stream Conference on Artificial Neural Networks and Expert Systems*, pp. 65–68. IEEE, 1993.
- Schaul, T., Quan, J., Antonoglou, I., and Silver, D. Prioritized experience replay. In *ICLR (Poster)*, 2016.
- Schwarzer, M., Anand, A., Goel, R., Hjelm, R. D., Courville, A., and Bachman, P. Data-efficient reinforcement learning with self-predictive representations. In *International Conference on Learning Representations*, 2020.
- Sharkey, N. E. and Sharkey, A. J. An analysis of catastrophic interference. *Connection Science*, 1995.
- Shteingart, H., Neiman, T., and Loewenstein, Y. The role of first impression in operant learning. *Journal of Experimental Psychology: General*, 142(2):476, 2013.
- Song, X., Jiang, Y., Tu, S., Du, Y., and Neyshabur, B. Observational overfitting in reinforcement learning. In *International Conference on Learning Representations*, 2019.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- Sutton, R. S. Learning to predict by the methods of temporal differences. *Machine learning*, 3(1):9–44, 1988.
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. MIT press, 2018.
- Taha, A., Shrivastava, A., and Davis, L. S. Knowledge evolution in neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12843–12852, 2021.
- Tassa, Y., Tunyasuvunakool, S., Muldal, A., Doron, Y., Liu, S., Bohez, S., Merel, J., Erez, T., Lillicrap, T., and Heess, N. dm_control: Software and tasks for continuous control, 2020.
- Teh, Y. W., Bapst, V., Czarnecki, W. M., Quan, J., Kirkpatrick, J., Hadsell, R., Heess, N., and Pascanu, R. Distal: Robust multitask reinforcement learning. In *NIPS*, 2017.
- Van Der Walt, S., Colbert, S. C., and Varoquaux, G. The numpy array: a structure for efficient numerical computation. *Computing in science & engineering*, 13(2):22–30, 2011.
- Van Hasselt, H., Guez, A., and Silver, D. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30, 2016.
- van Hasselt, H. P., Hessel, M., and Aslanides, J. When to use parametric models in reinforcement learning? *Advances in Neural Information Processing Systems*, 32:14322–14333, 2019.
- Van Rossum, G. and Drake Jr, F. L. *Python tutorial*, volume 620. Centrum voor Wiskunde en Informatica Amsterdam, 1995.
- Vanseijen, H. and Sutton, R. A deeper look at planning as learning from replay. In Bach, F. and Blei, D. (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 2314–2322, Lille, France, 07–09 Jul 2015. PMLR. URL <https://proceedings.mlr.press/v37/vanseijen15.html>.
- Wang, C., Wu, Y., Vuong, Q., and Ross, K. Striving for simplicity and performance in off-policy drl: Output normalization and non-uniform sampling. In *International Conference on Machine Learning*, pp. 10070–10080. PMLR, 2020.
- Yalnizyan-Carson, A. and Richards, B. A. Forgetting enhances episodic control with structured memories. *bioRxiv*, 2021.
- Yarats, D., Zhang, A., Kostrikov, I., Amos, B., Pineau, J., and Fergus, R. Improving sample efficiency in model-free reinforcement learning from images. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 10674–10681, 2021.
- Zhang, C., Vinyals, O., Munos, R., and Bengio, S. A study on overfitting in deep reinforcement learning. *arXiv preprint arXiv:1804.06893*, 2018.
- Zhang, C., Bengio, S., and Singer, Y. Are all layers created equal? *arXiv preprint arXiv:1902.01996*, 2019.
- Zhou, H., Vani, A., Larochelle, H., and Courville, A. Fortuitous forgetting in connectionist networks. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=ei3SY1_zYsE.

Task	Steps
walker-run	1×10^6
cheetah-run	1×10^6
acrobot-swingup	1×10^6
finger-turn_hard	1×10^6
fish-swim	1×10^6
humanoid-stand	1×10^6
humanoid-run	1×10^6
quadruped-run	1×10^6
swimmer-swimmer15	1×10^6
hopper-hop	1×10^6

Table 3. Tasks for SAC experiments and a number of training steps. Many of DMC tasks are solved by SAC in a matter of several thousand steps; we chose environments where SAC requires a substantial amount of training according to the reported results from https://github.com/denisyarats/pytorch_sac#results.

A. Experimental Details

We use an open-source JAX implementation (Kostrikov, 2021) of the SAC and DrQ algorithms and an open-source JAX implementation² of SPR. This SPR implementation exhibit a slightly higher aggregate performance than the scores in Schwarzer et al. (2020) based on a PyTorch implementation. All algorithms use default hyperparameters unless specified otherwise (for example, in experiments with replay ratio and n -step targets). SAC and DrQ experiments use 10 random seeds for evaluating the performance, SPR uses 20 random seeds.

Tables 3 and 4 report the sets of DeepMind Control Suite tasks for testing SAC and DrQ algorithms respectively. Note that SAC learns from dense states, while DrQ learns from raw pixel observations. SPR trains on a standard set of 26 tasks from the Atari 100k benchmark used by Kaiser et al. (2019); van Hasselt et al. (2019); Schwarzer et al. (2020). The list of all Atari environments is available in Table 7.

B. Ablations

Section 5.4 outlines the interaction of different moving parts of RL algorithms and our proposed reset strategy. This appendix elaborates on our observations and provides supporting figures to help understanding the effect of resets. We now show, one by one, how resetting behaves under a diverse set of controlled settings.

Replay buffer The reset mechanism we propose is a form of forgetting based on retaining all of the collected knowledge, stored in the replay buffer, and not retaining a part of the learned behavior, stored in the parameters of an agent’s

²https://github.com/MaxASchwarzer/dopamine/tree/ataril100k_spr

Task	Steps
walker-stand	5×10^5
finger-spin	5×10^5
cartpole-balance	5×10^5
cartpole-swingup	5×10^5
walker-walk	5×10^5
cartpole-balance_sparse	5×10^5
pendulum-swingup	1×10^6
hopper-stand	1×10^6
quadruped-walk	2×10^6
walker-run	2×10^6
finger-turn_easy	2×10^6
cheetah-run	2×10^6
acrobot-swingup	2×10^6
finger-turn_hard	2×10^6
cartpole-swingup_sparse	2×10^6
quadruped-run	2×10^6
reacher-easy	2×10^6
reacher-hard	2×10^6
hopper-hop	2×10^6

Table 4. Tasks for DrQ experiments and a number of training steps.

function approximators. How critical is it to preserve the knowledge in the replay buffer? We test its importance in DrQ by periodically resetting the replay buffer in addition to the last layers. Figure 8 shows that keeping the buffer is essential; resetting it amounts to learning almost from scratch. These results suggests that knowledge retention is significantly more important than behavior retention for preventing the negative effects of the primacy bias and simultaneously being able to recover from resets.

Initialization After each reset, the new parameters are sampled from a canonical initialization distribution followed by the algorithms. To understand whether the susceptibility of an agent to the primacy bias is just a consequence of an unlucky initialization of one of the layers of its neural networks, we run DrQ with resets, but set the value of the re-initialized parameters to the one they had at the beginning of training (i.e., by initializing them with the same seed). The results in Figure 8 show that the performance of this variant of our reset strategy is almost identical to the original version: the problem alleviated by resets is not one of pathological initializations, in line with findings of other works (Bjorck et al., 2021), but instead one resulting from the peculiar interactions happening when learning from a growing dataset of interactions.

Optimizer state As highlighted in the main text, resetting the optimizer’s moment estimates together with the corresponding neural network parameters have almost no effect. We show results demonstrating this for DrQ (see Figure 9). We believe this is mostly due to the momentum-based optimizers: with Adam with default parameters, the




Figure 8. Performance of the DrQ agent with standard resets, with same-seed resets, and with both buffer and last layer resetting (10 resets during training) on four DMC tasks. Resetting the replay buffer in addition to the last layers nullifies the learning progress, while preserving the random seed for drawing re-initialized parameters delivers almost the same results as standard resets.

first moment ($\beta_1 = 0.9$) will vanish in about 10 updates, the second moment ($\beta_2 = 0.999$) after 1000 updates. On the scale of our tasks, it is indeed a quite rapid recovery time.

Reset depth One of the two hyperparameters introduced by our reset strategy on top of any backbone algorithm is the number of layers of the agent’s neural networks to be re-initialized. We investigated the impact of this choice for both SPR and DrQ, while sticking for SAC to the default choice of re-initializing all networks. Results for DrQ in Figure 10 demonstrate that resetting the last layer yields slightly inferior performance to resetting the entire Q-learning head (3 layers). For SPR, as shown in Figure 13, we found the reverse to be true. Thus, how many layers to reset is a hyperparameter that may need tuning, and the choice can be informed by the difficulty of representation learning for each domain. We recommend starting the exploration of this hyperparameter from resetting the last 1-3 layers.

Which networks to reset In our experiments, we reset a subset of the value function parameters in SPR and a subset of the parameters of all the trained neural networks in DrQ and SAC. The latter two algorithms use three groups of function approximators: an actor, a critic, and a target critic. We investigate the impact of resetting each one of these modules in DrQ. Results in Figure 11 show that a simultaneous reset of all the neural networks is generally the most robust technique to improve performance over the backbone algorithm, while resetting the critic had the most impact on the performance. We have also tried a version of resets where

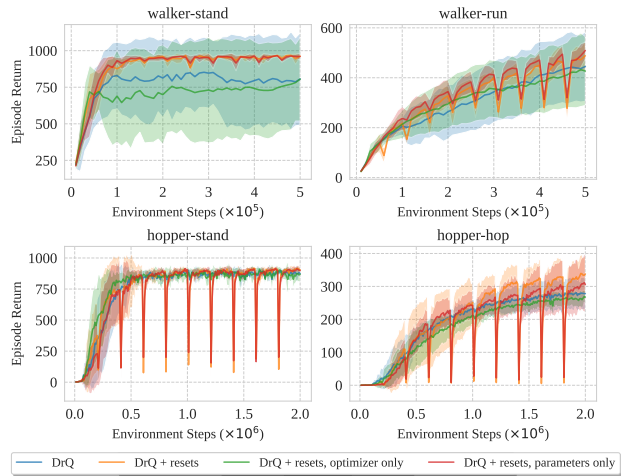


Figure 9. Performance of the DrQ agent with standard resets, with optimizer-only resets, and with parameter-only resets (10 resets during training) on four DMC tasks. Resetting the optimizer statistics does not alter training if the weights are preserved, while keeping the optimizer for re-initialized parameters delivers almost the same results as standard resets.

each weight is re-initialized with probability 0.5. Figure 10 shows that such a random subnetwork resetting was either on par or worse than the standard scheme.

Number of resets Intuitively, the primacy bias affects the agent in a progressively milder way after each reset. It is natural to ask whether a limited number of resets, or even a single one, is sufficient to overcome the effects of overfitting to initial experiences. We test this hypothesis using DrQ, showing in Figure 12 that, despite the first reset contributing the most to mitigating the primacy bias, it is not always sufficient to reach the same performance of the standard continual resetting strategy. As a default choice, we recommend using the reset periodicity resulting in 3-10 resets over the course of training.

Other regularizers Resets can be seen as a form of regularization because they implicitly constrain the final solutions to be not too far from the initial parameters. However, they specifically tackle the primacy bias better than other common forms of regularization. To test this conjecture, we repeat the heavy priming experiment on the quadruped-run task with standard L2 regularization of both critic and actor weights of SAC. We find that no value of regularization coefficient among the set $[10^{-5}, 3 \cdot 10^{-5}, 10^{-4}, 3 \cdot 10^{-4}, 10^{-3}, 3 \cdot 10^{-3}, 10^{-2}, 3 \cdot 10^{-2}, 10^{-1}, 3 \cdot 10^{-1}, 1.0]$ can overcome heavy priming, obtaining results almost identical to the ones reported on Figure 1. The heavy priming setting artificially creates the conditions for the effect of the primacy bias to be partic-

The Primacy Bias in Deep Reinforcement Learning

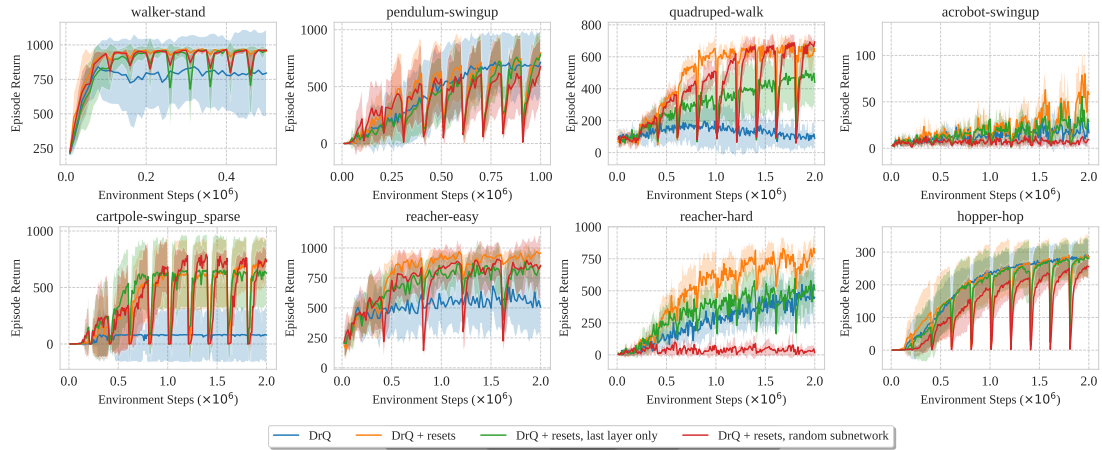


Figure 10. Performance of DrQ with resetting the 3-layer heads (i.e. standard resets), only the last layers, and random subnetworks. The overall performance of the latter two versions is either comparable or worse.




Figure 11. Performance of DrQ with resets of actor, critic, and target critic networks simultaneously (i.e. standard resets) and individually. Resetting the critic yields the predominant effect in most environments, but resetting all networks proved to be the most robust option.




Figure 12. Performance of DrQ when using a limited number of resets. The number of resets for reaching the best performance varies: in some environments, a single reset suffices to overcome the primacy bias, in other environments, keeping resetting continually is required.




Figure 13. Performance of SPR when resetting different groups of parameters. The right plot visualizes the percentage of improvement gained by resetting a certain group of parameters compared to no resets at all. Resetting the last layer only delivers a slightly higher IQM than resetting the 2-layer head, while resetting the whole network severely damages the performance.

ularly highlighted. To test whether resets offer superior performance also in the context of standard training of reinforcement learning algorithms, we compare the performance of SAC and DrQ enriched with standard regularization methods to that of SAC and DrQ augmented with resets. In particular, we leverage L2 regularization of both the actor and the critic, as well as dropout (Srivastava et al., 2014). We report in Table 5 the best value over the grid $[10^{-4}, 5 \cdot 10^{-4}, 10^{-3}, 5 \cdot 10^{-3}]$ suggested by (Liu et al., 2021) for L2 and the standard grid $[0.5, 0.1]$ for dropout. For SAC-based approaches, we also sweep over the replay ratios $[1, 9, 32]$ and report the best result for other regularizers. The table shows that not only standard regularization methods are not better than resets in the context of these RL algorithms, but that they do not provide any benefit, on aggregate performance, compared to the baselines.

Method	IQM	Median	Mean
SAC	501 (389, 609)	475 (407, 563)	484 (420, 548)
SAC + resets	656 (549, 753)	617 (538, 681)	607 (547, 667)
SAC + dropout	219 (160, 285)	254 (204, 307)	258 (216, 300)
SAC + L2	412 (299, 524)	415 (337, 495)	416 (351, 481)
DrQ	569 (475, 662)	521 (470, 600)	535 (481, 589)
DrQ + resets	762 (704, 815)	680 (625, 731)	677 (632, 720)
DrQ + dropout	492 (414, 567)	480 (420, 541)	479 (431, 527)
DrQ + L2	463 (362, 566)	473 (403, 541)	472 (415, 529)

Table 5. Comparison of the performance of SAC and DrQ when augmented with standard regularization techniques and resets. We leverage 10 runs and the same set of evaluation tasks reported in Table 3 and Table 4.

C. Per-Environment and Additional Results

The remainder of the appendix presents results for each task and supplementary plots for training with varying replay ratios and n -step targets.

Figure 14 demonstrates learning curves for SPR. We note that the low loss and high parameter norm for high n and replay ratios might indicate the symptoms of overfitting. Whilst resets implicitly control the weight norm, doing so explicitly through L2 regularization proved to be less effec-

tive for mitigating heavy priming.

Table 6 presents the aggregate metrics for the combinations of n and replay ratios in SAC. We additionally probe extreme replay ratios of 128 and 256 and observe that, even in these cases, learning with resets delivers meaningful performance, while the no-reset agent achieves near-zero returns.

Lastly, per-environment training curves for SAC as well as the results for varying replay ratios and n -step targets are available in Figures 15 and 17 respectively. Per-environment training curves for DrQ are available in Figure 18. Table 7 provides scores for SPR in all Atari 100k tasks.

The Primacy Bias in Deep Reinforcement Learning

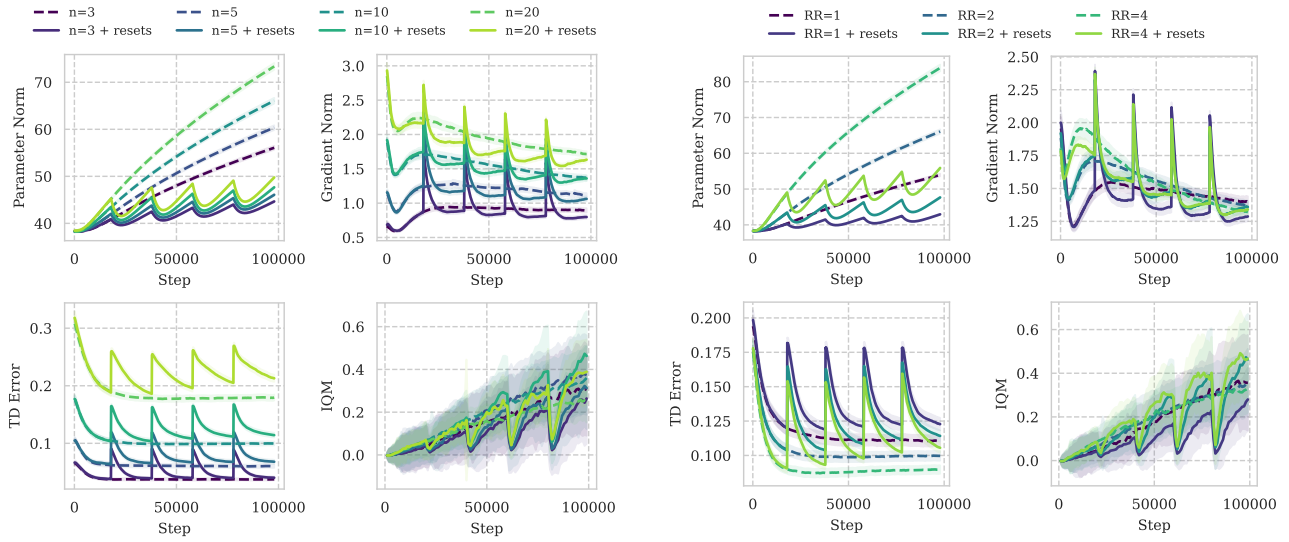


Figure 14. Training curves for SPR with resets on Atari 100k with different n -step targets (left) and varying replay ratios (right). Resets increase TD errors and temporarily increase gradient norms for all values of n and RR, while implicitly regularizing parameter norms.

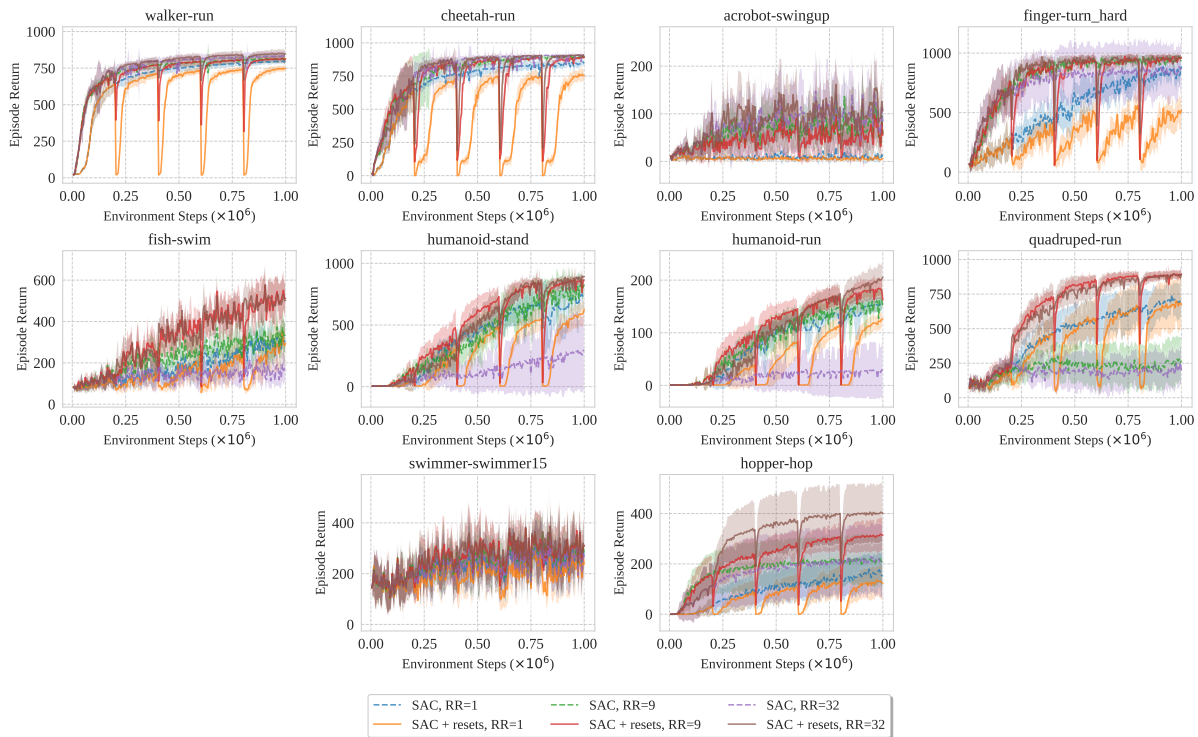


Figure 15. Per-environment training curves for SAC for various replay ratio (RR) values and a fixed value of $n = 1$.

The Primacy Bias in Deep Reinforcement Learning

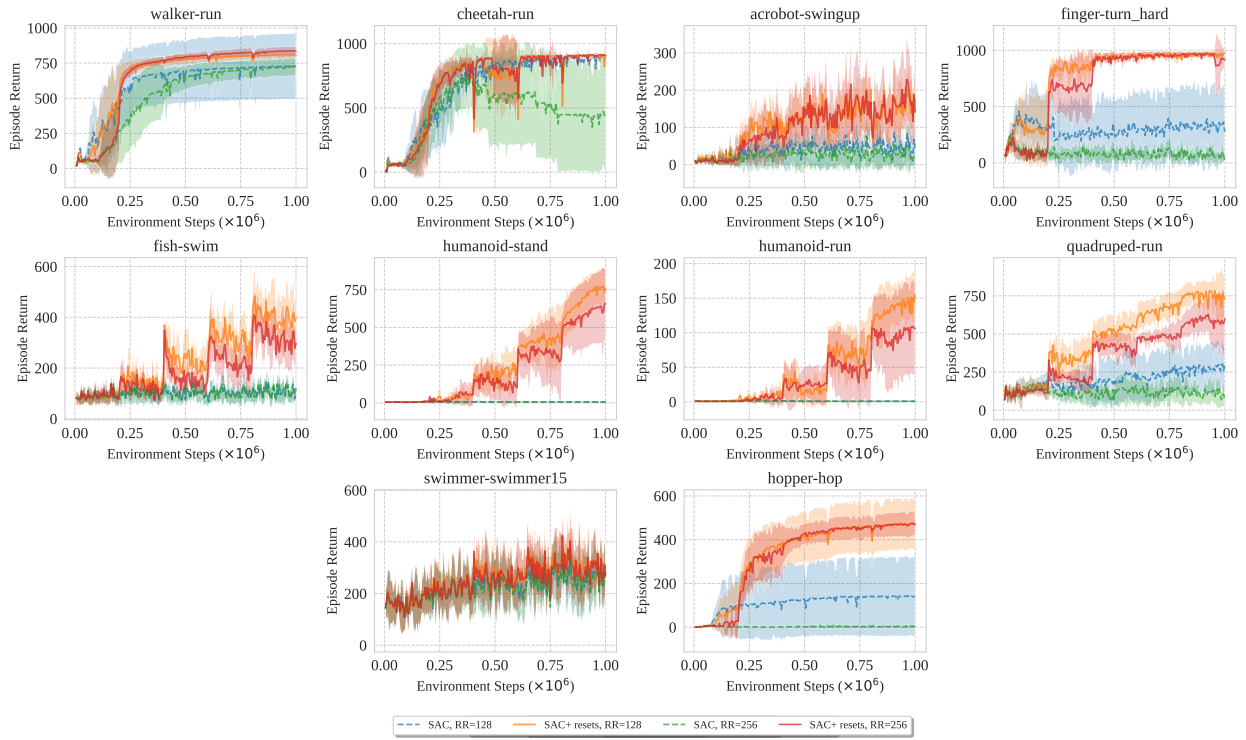


Figure 16. Per-environment training curves for SAC for the extreme replay ratio (RR) values of 128 and 256 and a fixed value of $n = 1$. While a standard learning algorithm struggles to make any progress, resets allow to achieve reasonable performance in this regime.

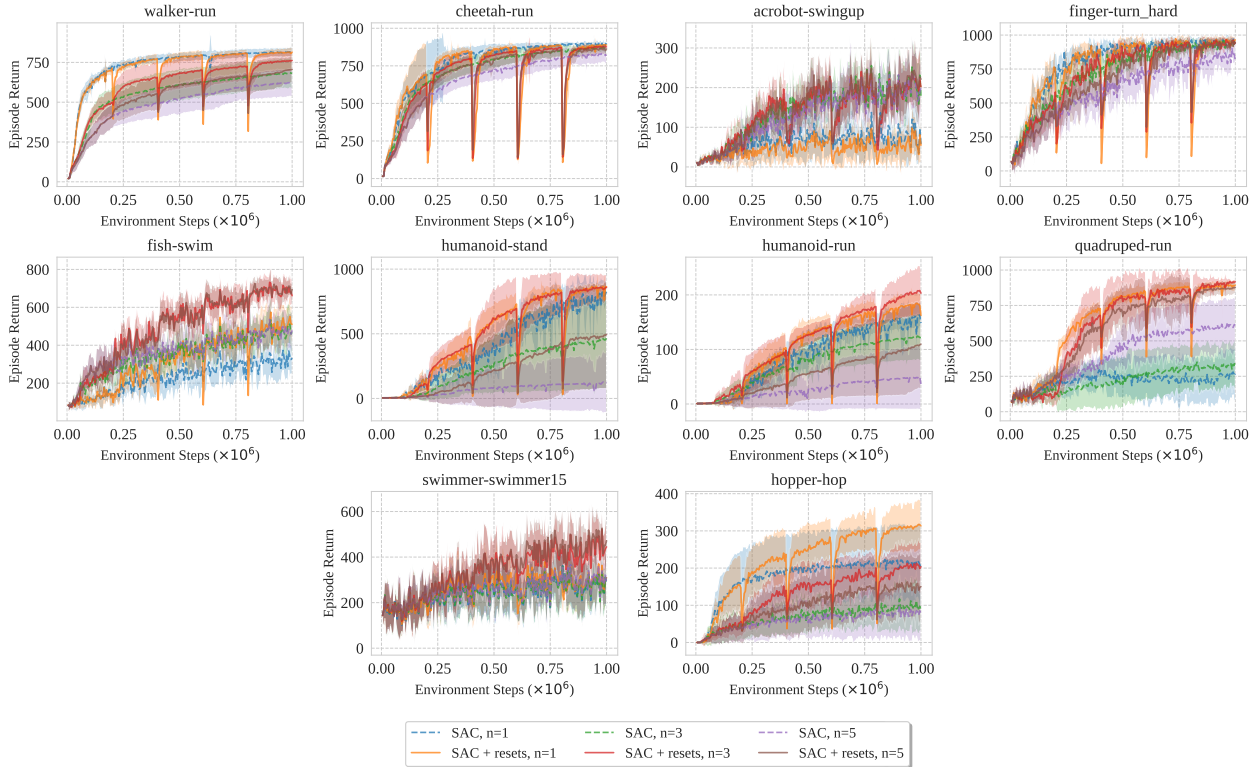


Figure 17. Per-environment training curves for SAC for various n and a fixed value of replay ratio $RR = 9$.

RR	$n=1$		$n=3$		$n=5$	
	resets	no resets	resets	no resets	resets	no resets
1	401 (308,497), 500 (389,609)		431 (366,499), 452 (347,556)		434 (333,533), 399 (304,493)	
9	628 (514,730), 459 (341,576)		656 (549,752), 413 (319,517)		597 (490,699), 392 (298,488)	
32	651 (542,750), 301 (193,430)		642 (541,737), 342 (241,452)		588 (482,684), 291 (200,394)	
128	584 (487,679), 149 (86,242)		—		—	
256	520 (418,619), 39 (18,79)		—		—	

(a) IQM

RR	$n=1$		$n=3$		$n=5$	
	resets	no resets	resets	no resets	resets	no resets
1	415 (342,478), 474 (406,562)		425 (382,481), 448 (375,527)		409 (346,492), 398 (328,475)	
9	575 (500,656), 473 (395,557)		616 (537,680), 436 (365,519)		547 (476,632), 403 (333,480)	
32	602 (527,677), 372 (297,469)		599 (532,674), 398 (317,476)		554 (475,626), 352 (279,433)	
128	568 (496,640), 268 (192,349)		—		—	
256	518 (444,595), 152 (92,221)		—		—	

(b) Median

RR	$n=1$		$n=3$		$n=5$	
	resets	no resets	resets	no resets	resets	no resets
1	410 (354,467), 484 (419,549)		433 (393,473), 451 (387,515)		419 (360,478), 407 (348,467)	
9	577 (511,642), 476 (409,543)		607 (547,666), 443 (381,504)		553 (488,617), 407 (346,469)	
32	600 (537,662), 384 (314,456)		601 (543,659), 397 (332,464)		549 (486,611), 358 (294,422)	
128	566 (507,626), 275 (212,341)		—		—	
256	520 (457,581), 165 (115,220)		—		—	

(c) Mean

Table 6. Full results for SAC in terms of IQM, median, and mean performance across tasks.

The Primacy Bias in Deep Reinforcement Learning




Figure 18. Per-environment training curves for DrQ with and without resets on DMC.

	Random	Human	DER	DrQ(ϵ)	SPR	SPR+resets
Alien	227.8	7127.7	802.3	865.2	901.2	911.2
Amidar	5.8	1719.5	125.9	137.8	225.4	201.7
Assault	222.4	742.0	561.5	579.6	658.6	953.0
Asterix	210.0	8503.3	535.4	763.6	1095.0	1005.8
Bank Heist	14.2	753.1	185.5	232.9	484.7	547.0
Battle Zone	2360.0	37187.5	8977.0	10165.3	10873.5	8821.2
Boxing	0.1	12.1	-0.3	9.0	27.6	32.2
Breakout	1.7	30.5	9.2	19.8	16.9	23.4
Chopper Command	811.0	7387.8	925.9	844.6	1454.0	1680.6
Crazy Climber	10780.5	35829.4	34508.6	21539.0	23596.9	28936.2
Demon Attack	152.1	1971.0	627.6	1321.5	1291.7	2778.0
Freeway	0.0	29.6	20.9	20.3	9.7	18.0
Frostbite	65.2	4334.7	871.0	1014.2	1746.2	1834.3
Gopher	257.6	2412.5	467.0	621.6	642.4	930.4
Hero	1027.0	30826.4	6226.0	4167.9	7554.5	6735.6
Jamesbond	29.0	302.8	275.7	349.1	383.2	415.7
Kangaroo	52.0	3035.0	581.7	1088.4	1674.8	2190.6
Krull	1598.0	2665.5	3256.9	4402.1	3412.1	4772.4
Kung Fu Master	258.5	22736.3	6580.1	11467.4	16688.6	14682.1
Ms Pacman	307.3	6951.6	1187.4	1218.1	1334.1	1324.6
Pong	-20.7	14.6	-9.7	-9.1	2.1	-9.0
Private Eye	24.9	69571.3	72.8	3.5	76.1	82.2
Qbert	163.9	13455.0	1773.5	1810.7	3816.2	3955.3
Road Runner	11.5	7845.0	11843.4	11211.4	13588.5	13088.2
Seaquest	68.4	42054.7	304.6	352.3	519.7	655.6
Up N Down	533.4	11693.2	3075.0	4324.5	8873.4	60185.0
Median HNS	0.000	1.000	0.189	0.313	0.453	0.512
Mean HNS	0.000	1.000	0.350	0.465	0.579	0.901
#Games > Human	0	0	2	3	4	7
#Games > 0	0	28	25	25	26	26

Table 7. Raw per-game scores and aggregate human-normalized scores (HNS) for SPR with resets and other methods on all 26 games in the Atari 100k benchmark. We report performance for SPR and SPR + resets from our codebase, averaged over 20 random seeds per game; other scores are taken from Agarwal et al. (2021) and use 100 random seeds.