

# PRIORITIZED EXPERIENCE REPLAY

Tom Schaul, John Quan, Ioannis Antonoglou and David Silver

Google DeepMind

{schaul, johnquan, ioannisa, davidsilver}@google.com

## ABSTRACT

Experience replay lets online reinforcement learning agents remember and reuse experiences from the past. In prior work, experience transitions were uniformly sampled from a replay memory. However, this approach simply replays transitions at the same frequency that they were originally experienced, regardless of their significance. In this paper we develop a framework for prioritizing experience, so as to replay important transitions more frequently, and therefore learn more efficiently. We use prioritized experience replay in Deep Q-Networks (DQN), a reinforcement learning algorithm that achieved human-level performance across many Atari games. DQN with prioritized experience replay achieves a new state-of-the-art, outperforming DQN with uniform replay on 42 out of 57 games.

## 1 INTRODUCTION

Online reinforcement learning (RL) agents incrementally update their parameters (of the policy, value function or model) while they observe a stream of experience. In their simplest form, they discard incoming data immediately, after a single update. Two issues with this are (a) strongly correlated updates that break the i.i.d. assumption of many popular stochastic gradient-based algorithms, and (b) the rapid forgetting of possibly rare experiences that would be useful later on.

*Experience replay* (Lin, 1992) addresses both of these issues: with experience stored in a replay memory, it becomes possible to break the temporal correlations by mixing more and less recent experience for the updates, and rare experience will be used for more than just a single update. This was demonstrated in the Deep Q-Network (DQN) algorithm (Mnih et al., 2013; 2015), which stabilized the training of a value function, represented by a deep neural network, by using experience replay. Specifically, DQN used a large sliding window replay memory, sampled from it uniformly at random, and effectively revisited each transition<sup>1</sup> eight times. In general, experience replay can reduce the amount of experience required to learn, and replace it with more computation and more memory – which are often cheaper resources than the RL agent’s interactions with its environment.

In this paper, we investigate how *prioritizing* which transitions are replayed can make experience replay more efficient and effective than if all transitions are replayed uniformly. The key idea is that an RL agent can learn more effectively from some transitions than from others. Transitions may be more or less surprising, redundant, or task-relevant. Some transitions may not be immediately useful to the agent, but might become so when the agent competence increases (Schmidhuber, 1991). Experience replay liberates online learning agents from processing transitions in the exact order they are experienced. Prioritized replay further liberates agents from considering transitions with the same frequency that they are experienced.

In particular, we propose to more frequently replay transitions with high expected learning progress, as measured by the magnitude of their temporal-difference (TD) error. This prioritization can lead to a loss of diversity, which we alleviate with stochastic prioritization, and introduce bias, which we correct with importance sampling. Our resulting algorithms are robust and scalable, which we demonstrate on the Atari 2600 benchmark suite, where we obtain faster learning and state-of-the-art performance.

<sup>1</sup> A transition is the atomic unit of interaction in RL, in our case a tuple of (state  $S_{t-1}$ , action  $A_{t-1}$ , reward  $R_t$ , discount  $\gamma_t$ , next state  $S_t$ ). We choose this for simplicity, but most of the arguments in this paper also hold for a coarser ways of chunking experience, e.g. into sequences or episodes.

## 2 BACKGROUND

Numerous neuroscience studies have identified evidence of experience replay in the hippocampus of rodents, suggesting that sequences of prior experience are replayed, either during awake resting or sleep. Sequences associated with rewards appear to be replayed more frequently (Atherton et al., 2015; Ólafsdóttir et al., 2015; Foster & Wilson, 2006). Furthermore, experiences with high magnitude of TD error appear to be replayed more often (Singer & Frank, 2009; McNamara et al., 2014).

It is well-known that planning algorithms such as value iteration can be made more efficient by prioritizing updates in an appropriate order. *Prioritized sweeping* (Moore & Atkeson, 1993; Andre et al., 1998) selects which state to update next, prioritized according to the change in value, if that update was executed. The TD error provides one way to measure these priorities (van Seijen & Sutton, 2013). Our approach uses a similar prioritization method, but for model-free RL rather than model-based planning. Furthermore, we use a stochastic prioritization that is more robust when learning a function approximator from samples.

TD-errors have also been used as a prioritization mechanism for determining where to focus resources, for example when choosing where to explore (White et al., 2014) or which features to select (Geramifard et al., 2011; Sun et al., 2011).

In supervised learning, there are numerous techniques to deal with imbalanced datasets when class identities are known, including re-sampling, under-sampling and over-sampling techniques, possibly combined with ensemble methods (Galar et al., 2012). A form of non-uniform sampling based on error, and with importance sampling correction was introduced by Hinton (2007), and led to a reported 3x speed-up on MNIST digit classification.

There have been several proposed methods for playing Atari with deep reinforcement learning, including deep Q-networks (DQN) (Mnih et al., 2013; 2015; Guo et al., 2014; Stadie et al., 2015; Nair et al., 2015; Bellemare et al., 2016), and the *Double DQN* algorithm (van Hasselt et al., 2016), which is the current published state-of-the-art. Simultaneously with our work, an architectural innovation that separates advantages from the value function (see the co-submission by Wang et al., 2015) has also led to substantial improvements on the Atari benchmark.

## 3 PRIORITIZED REPLAY

Using a replay memory leads to design choices at two levels: which experiences to store, and which experiences to replay (and how to do so). This paper addresses only the latter: making the most effective use of the replay memory for learning, assuming that its contents are outside of our control (but see also Section 6).

### 3.1 A MOTIVATING EXAMPLE

To understand the potential gains of prioritization, we introduce an artificial ‘Blind Cliffwalk’ environment (described in Figure 1, left) that exemplifies the challenge of exploration when rewards are rare. With only  $n$  states, the environment requires an exponential number of random steps until the first non-zero reward; to be precise, the chance that a random sequence of actions will lead to the reward is  $2^{-n}$ . Furthermore, the most relevant transitions (from rare successes) are hidden in a mass of highly redundant failure cases (similar to a bipedal robot falling over repeatedly, before it discovers how to walk).

We use this example to highlight the difference between the learning times of two agents. Both agents perform Q-learning updates on transitions drawn from the same replay memory. The first agent replays transitions uniformly at random, while the second agent invokes an oracle to prioritize transitions. This oracle greedily selects the transition that maximally reduces the global loss in its current state (in hindsight, after the parameter update). For the details of the setup, see Appendix B.1. Figure 1 (right) shows that picking the transitions in a good order can lead to exponential speed-ups over uniform choice. Such an oracle is of course not realistic, yet the large gap motivates our search for a practical approach that improves on uniform random replay.

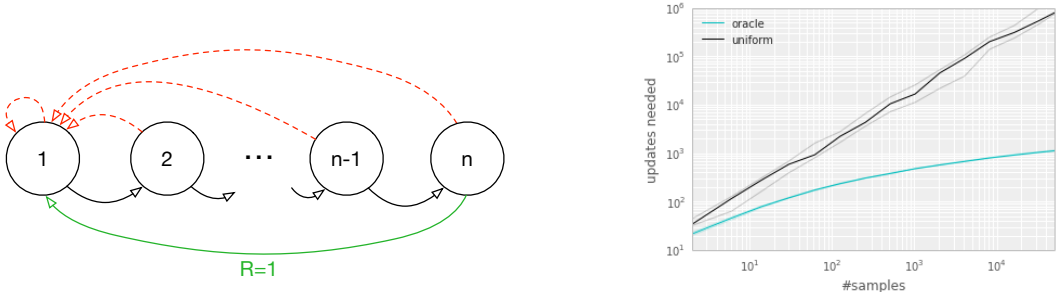


Figure 1: **Left:** Illustration of the ‘Blind Cliffwalk’ example domain: there are two actions, a ‘right’ and a ‘wrong’ one, and the episode is terminated whenever the agent takes the ‘wrong’ action (dashed red arrows). Taking the ‘right’ action progresses through a sequence of  $n$  states (black arrows), at the end of which there is a final reward of 1 (green arrow); reward is 0 everywhere else. We choose the representation such that generalization over what action is ‘right’ is impossible. **Right:** Median number of learning steps required to learn the value function as a function of the size of the total number of transitions in the replay memory. Note the log-log scale, which highlights the exponential speed-up from replaying with an oracle (bright blue), compared to uniform replay (black); faint lines are min/max values from 10 independent runs.

### 3.2 PRIORITIZING WITH TD-ERROR

The central component of prioritized replay is the criterion by which the importance of each transition is measured. One idealised criterion would be the amount the RL agent can expect to learn from that transition in its current state (expected learning progress). While this measure is not directly accessible, a reasonable proxy is the magnitude of a transition’s TD error  $\delta$ , which indicates how ‘surprising’ or unexpected the transition is: specifically, how far the value is from its next-step bootstrap estimate (Andre et al., 1998). This is particularly suitable for incremental, online RL algorithms, such as SARSA or Q-learning, that already compute the TD-error and update the parameters in proportion to  $\delta$ . While appropriate in many circumstances, the TD-error can sometimes be a poor estimate as well; see Appendix A for a discussion of alternatives.

To demonstrate the potential effectiveness of prioritizing replay by TD error, we compare the uniform and oracle baselines in the Blind Cliffwalk to a ‘greedy TD-error prioritization’ algorithm. This algorithm stores the last encountered TD error along with each transition in the replay memory. The transition with the largest absolute TD error is replayed from the memory. A Q-learning update is applied to this transition, which updates weights in proportion to the TD error. This new TD error replaces the previous value for that transition. New transitions arrive without a known TD-error, so we put them at maximal priority in order to guarantee that all experience is seen at least once. Figure 2 (left), shows that this algorithm results in a substantial reduction in the length of time required to solve the Blind Cliffwalk, coming close in performance to the oracle prioritisation.<sup>2</sup>

**Implementation:** To scale to large memory sizes  $N$ , we use a binary heap data structure for the priority queue, for which finding the maximum priority transition when sampling is  $O(1)$ , and updating priorities (with the new TD-error after a learning step) is  $O(\log N)$ . See Appendix B.2.1 for more details.

### 3.3 STOCHASTIC PRIORITIZATION

However, greedy TD-error prioritization has several issues. First, to avoid expensive sweeps over the entire replay memory, TD errors are only updated for the transitions that are replayed. One consequence is that transitions that have a low TD error on first visit may not be replayed for a long time (which effectively means never with a sliding window replay memory). Further, it is sensitive to noise spikes (e.g. when rewards are stochastic), which can be exacerbated by bootstrapping, where approximation errors appear as another source of noise. Finally, greedy prioritization focuses on a small subset of the experience: errors shrink slowly, especially when using function approximation,

<sup>2</sup> Note that a random (or optimistic) initialization of the Q-values is necessary with greedy prioritization. If initializing with zero instead, unrewarded transitions would appear to have zero error initially, be placed at the bottom of the queue, and not be revisited until the error on other transitions drops below numerical precision.

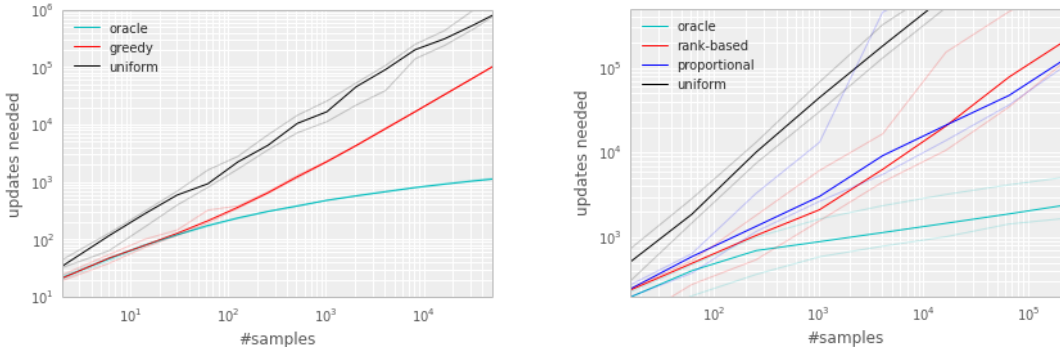


Figure 2: Median number of updates required for Q-learning to learn the value function on the Blind Cliffwalk example, as a function of the total number of transitions (only a single one of which was successful and saw the non-zero reward). Faint lines are min/max values from 10 random initializations. Black is uniform random replay, cyan uses the hindsight-oracle to select transitions, red and blue use prioritized replay (rank-based and proportional respectively). The results differ by multiple orders of magnitude, thus the need for a log-log plot. In both subplots it is evident that replaying experience in the right order makes an enormous difference to the number of updates required. See Appendix B.1 for details. **Left:** Tabular representation, greedy prioritization. **Right:** Linear function approximation, both variants of stochastic prioritization.

meaning that the initially high error transitions get replayed frequently. This lack of diversity that makes the system highly prone to over-fitting.

To overcome these issues, we introduce a stochastic sampling method that interpolates between pure greedy prioritization and uniform random sampling. We ensure that the probability of being sampled is monotonic in a transition’s priority, while guaranteeing a non-zero probability even for the lowest-priority transition. Concretely, we define the probability of sampling transition  $i$  as

$$P(i) = \frac{p_i^\alpha}{\sum_k p_k^\alpha} \quad (1)$$

where  $p_i > 0$  is the priority of transition  $i$ . The exponent  $\alpha$  determines how much prioritization is used, with  $\alpha = 0$  corresponding to the uniform case.

The first variant we consider is the direct, proportional prioritization where  $p_i = |\delta_i| + \epsilon$ , where  $\epsilon$  is a small positive constant that prevents the edge-case of transitions not ever being revisited once their error is zero. The second variant is an indirect, rank-based prioritization where  $p_i = \frac{1}{\text{rank}(i)}$ , where  $\text{rank}(i)$  is the rank of transition  $i$  when the replay memory is sorted according to  $|\delta_i|$ . In this case,  $P$  becomes a power-law distribution with exponent  $\alpha$ . Both distributions are monotonic in  $|\delta|$ , but the latter is more robust, as it is not sensitive to outliers. Both variants of stochastic prioritization lead to large speed-ups over the uniform baseline on the Cliffwalk task, as shown on Figure 2 (right).

**Implementation:** To efficiently sample from distribution (1), the complexity cannot depend on  $N$ . For the rank-based variant, we can approximate the cumulative density function with a piecewise linear function with  $k$  segments of equal probability. The segment boundaries can be precomputed (they change only when  $N$  or  $\alpha$  change). At runtime, we sample a segment, and then sample uniformly among the transitions within it. This works particularly well in conjunction with a minibatch-based learning algorithm: choose  $k$  to be the size of the minibatch, and sample exactly one transition from each segment – this is a form of stratified sampling that has the added advantage of balancing out the minibatch (there will always be exactly one transition with high magnitude  $\delta$ , one with medium magnitude, etc). The proportional variant is different, but it also admits an implementation that can be efficiently updated and sampled from. This implementation uses a ‘sum-tree’ data structure, where every node is the sum of its children, with the priorities as the leaf nodes. See Appendix B.2.1 for more additional details.

### 3.4 ANNEALING THE BIAS

The estimation of the expected value with stochastic updates relies on those updates corresponding to the same distribution as its expectation. Prioritized replay introduces bias because it changes this

**Algorithm 1** Double DQN with proportional prioritization

---

```

1: Input: minibatch  $k$ , step-size  $\eta$ , replay period  $K$  and size  $N$ , exponents  $\alpha$  and  $\beta$ , budget  $T$ .
2: Initialize replay memory  $\mathcal{H} = \emptyset$ ,  $\Delta = 0$ ,  $p_1 = 1$ 
3: Observe  $S_0$  and choose  $A_0 \sim \pi_\theta(S_0)$ 
4: for  $t = 1$  to  $T$  do
5:   Observe  $S_t, R_t, \gamma_t$ 
6:   Store transition  $(S_{t-1}, A_{t-1}, R_t, \gamma_t, S_t)$  in  $\mathcal{H}$  with maximal priority  $p_t = \max_{i < t} p_i$ 
7:   if  $t \equiv 0 \pmod K$  then
8:     for  $j = 1$  to  $k$  do
9:       Sample transition  $j \sim P(j) = p_j^\alpha / \sum_i p_i^\alpha$ 
10:      Compute importance-sampling weight  $w_j = (N \cdot P(j))^{-\beta} / \max_i w_i$ 
11:      Compute TD-error  $\delta_j = R_j + \gamma_j Q_{\text{target}}(S_j, \arg \max_a Q(S_j, a)) - Q(S_{j-1}, A_{j-1})$ 
12:      Update transition priority  $p_j \leftarrow |\delta_j|$ 
13:      Accumulate weight-change  $\Delta \leftarrow \Delta + w_j \cdot \delta_j \cdot \nabla_\theta Q(S_{j-1}, A_{j-1})$ 
14:    end for
15:    Update weights  $\theta \leftarrow \theta + \eta \cdot \Delta$ , reset  $\Delta = 0$ 
16:    From time to time copy weights into target network  $\theta_{\text{target}} \leftarrow \theta$ 
17:  end if
18:  Choose action  $A_t \sim \pi_\theta(S_t)$ 
19: end for

```

---

distribution in an uncontrolled fashion, and therefore changes the solution that the estimates will converge to (even if the policy and state distribution are fixed). We can correct this bias by using importance-sampling (IS) weights

$$w_i = \left( \frac{1}{N} \cdot \frac{1}{P(i)} \right)^\beta$$

that fully compensates for the non-uniform probabilities  $P(i)$  if  $\beta = 1$ . These weights can be folded into the Q-learning update by using  $w_i \delta_i$  instead of  $\delta_i$  (this is thus *weighted* IS, not ordinary IS, see e.g. Mahmood et al., 2014). For stability reasons, we always normalize weights by  $1 / \max_i w_i$  so that they only scale the update downwards.

In typical reinforcement learning scenarios, the unbiased nature of the updates is most important near convergence at the end of training, as the process is highly non-stationary anyway, due to changing policies, state distributions and bootstrap targets; we hypothesize that a small bias can be ignored in this context (see also Figure 11 in the appendix for a case study of full IS correction on Atari). We therefore exploit the flexibility of *annealing* the amount of importance-sampling correction over time, by defining a schedule on the exponent  $\beta$  that reaches 1 only at the end of learning. In practice, we linearly anneal  $\beta$  from its initial value  $\beta_0$  to 1. Note that the choice of this hyper-parameter interacts with choice of prioritization exponent  $\alpha$ ; increasing both simultaneously makes prioritization more aggressive at the same time as correcting for it more strongly.

Importance sampling has another benefit when combined with prioritized replay in the context of non-linear function approximation (e.g. deep neural networks): the gradient magnitudes (and thus the effective steps in parameter space) are proportional to the error, however, because in the non-linear case the 1st-order approximation of the gradient is only reliable locally, this means that large steps can be very disruptive, and the global step-size must be small to prevent them. Instead, doing many small updates on high-error transitions also alleviates the issue because the Taylor expansion is re-approximated, allowing the algorithm follow the curvature of highly non-linear optimization landscapes. This is exactly what happens in our approach: prioritization makes sure the high-error transition is seen many times, and the IS correction weight makes each resulting step a small one.

We combine our prioritized replay algorithm into a full-scale reinforcement learning agent, based on the state-of-the-art Double DQN algorithm. Our principal modification is to replace the uniform random sampling used by Double DQN with our stochastic prioritization and importance sampling methods (see Algorithm 1).

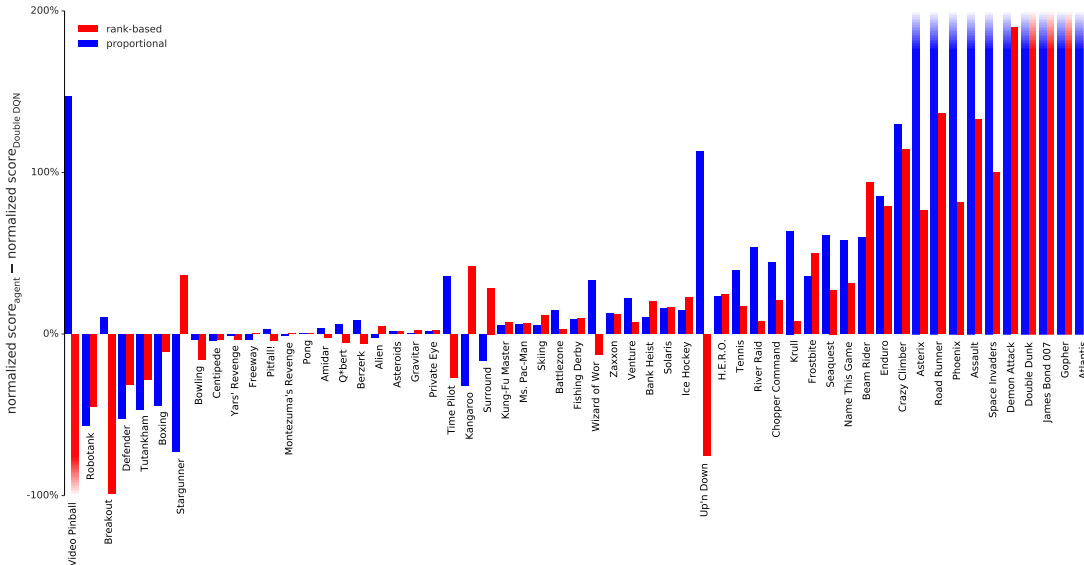


Figure 3: Difference in normalized score (the gap between random and human is 100%) on 57 games with human starts, comparing Double DQN with and without prioritized replay (rank-based variant in red, proportional in blue), showing substantial improvements in most games. Exact scores are in Table 6. See also Figure 8 where regular DQN is the baseline.

#### 4 ATARI EXPERIMENTS

With all these concepts in place, we now investigate to what extent replay with such prioritized sampling can improve performance in realistic problem domains. For this, we chose the collection of Atari benchmarks (Bellemare et al., 2012) with their end-to-end RL from vision setup, because they are popular and contain diverse sets of challenges, including delayed credit assignment, partial observability, and difficult function approximation (Mnih et al., 2015; van Hasselt et al., 2016). Our hypothesis is that prioritized replay is generally useful, so that it will make learning with experience replay more efficient without requiring careful problem-specific hyper-parameter tuning.

We consider two baseline algorithms that use uniform experience replay, namely the version of the DQN algorithm from the Nature paper (Mnih et al., 2015), and its recent extension Double DQN (van Hasselt et al., 2016) that substantially improved the state-of-the-art by reducing the over-estimation bias with Double Q-learning (van Hasselt, 2010). We use the tuned version of the Double DQN algorithm. The most relevant component for this work is how the replay memory functions: all experienced transitions are stored in a sliding window memory that retains the last  $10^6$  transitions and automatically drops older ones. The learning algorithm processes minibatches of 32 transitions for each update, which are sampled uniformly from the memory. There is one minibatch update for each 4 new transitions entering the memory, so all experience is replayed 8 times on average. Rewards and TD-errors are clipped to fall within  $[-1, 1]$  for stability reasons.

We use the identical neural network architecture, learning algorithm, replay memory and evaluation setup as for the baselines (see Appendix B.2). The only difference is the mechanism for sampling transitions from the replay memory, with is now done according to Algorithm 1 instead of uniformly. We compare the baselines to both variants of prioritized replay (rank-based and proportional prioritization).

Only a single hyper-parameter adjustment was necessary compared to the baseline: Given that prioritized replay picks high-error transitions more often, the typical gradient magnitudes are larger: we thus had reduced the step-size  $\eta$  by a factor 4 compared to the (Double) DQN setup. For the  $\alpha$  and  $\beta_0$  hyper-parameters that are introduced by prioritization, we did a coarse grid search (evaluated on a subset of 8 games), and found the sweet spot to be  $\alpha = 0.7$ ,  $\beta_0 = 0.5$  for the rank-based variant and  $\alpha = 0.6$ ,  $\beta_0 = 0.4$  for the proportional variant. These choices are trading off aggressiveness with robustness, by reducing  $\alpha$  and/or increasing  $\beta$  and it is always possible to revert to a behavior closer to the baseline.

We produce the results by running each variant with a single hyper-parameter setting across all games, as was done for the baselines. Our main evaluation metric is the policy quality in terms of average score per episode, given start states sampled from human traces (as introduced in Nair et al., 2015 and used in van Hasselt et al., 2016), which requires more robustness and generalization (the agent cannot rely on repeating a single memorized trajectory). These results are summarized in Table 6, raw scores and additional evaluation metrics can be found in Tables 8 and 9 in the Appendix.

We find that adding prioritized replay to DQN leads to a substantial improvement in score on 42 out of 57 games (compare columns 3 and 4 of Table 6 or Figure 8 in the appendix), with the median normalized performance score across 57 games increased from 69% to 97%. Furthermore, we find that the boost from prioritized experience replay is *complementary* to the one from introducing Double Q-learning into DQN: performance increases another notch, leading to the current state-of-the-art on the Atari benchmark (see Figure 3). Compared to Double DQN, the median performance increased from 110% to 128%, and the mean performance from 389% to 551% bringing additional games such as River Raid, Seaquest and Surround to a human level for the first time, and making large jumps on others (e.g. Atlantis, Gopher, James Bond 007 or Space Invaders). Note that mean performance is not a very reliable metric because a single game (Video Pinball) has a dominant contribution. Adding in prioritized replay gives a performance boost on almost all games, and learning is also faster (see Figures 6 and 7). The learning curves on Figure 6 illustrate that while the two variants of prioritization usually lead to similar results, there are games where one of them remains close to the Double DQN baseline while the other one leads to a big boost, for example Double Dunk or Surround for the rank-based variant, and Alien, Asterix, Enduro, Phoenix or Space Invaders for the proportional variant. Another observation from the learning curves is that compared to the uniform baseline, prioritization is effective at reducing the delay until performance gets off the ground in games that suffer from it, such as Battlezone, Zaxxon or Frostbite – but it fails to solve Montezuma’s Revenge.

	Evaluation method	DQN		Double DQN		
		baseline	rank-based	baseline	rank-based	proportional
<b>Median</b>	human starts	69%	97%	110%	113%	<b>128%</b>
	no-ops	79%	115%	118%	124%	<b>140%</b>
<b>Mean</b>	human starts	242%	311%	389%	454%	<b>551%</b>
	no-ops	433%	541%	641%	736%	<b>912%</b>
<b>&gt; baseline</b>	human starts	–	42	–	41	<b>44</b>
	no-ops	–	40	–	40	<b>43</b>
<b>&gt; human</b>	human starts	21	27	30	<b>33</b>	<b>33</b>
	no-ops	24	31	33	<b>39</b>	<b>37</b>

Table 1: Summary of normalized scores. See Tables 6 and 7 in the appendix for full results.

## 5 DISCUSSION

In the head-to-head comparison between rank-based prioritization and proportional prioritization, we expected the rank-based variant to be much more robust, because it is not affected by outliers or by absolute magnitudes of errors. Furthermore, its heavy-tail property also guarantees that samples will be diverse, and the stratified sampling from partitions of different error magnitudes will keep the total minibatch gradient at a stable magnitude throughout training. On the other hand, the ranks make the algorithm blind to the relative error scales, which could incur a performance drop when there is structure in the distribution of errors to be exploited, such as in sparse reward scenarios. Perhaps surprisingly, both variants perform similarly in practice; we suspect this is due to the heavy use of clipping (of rewards and TD-errors) in the DQN algorithm, which protects it from outliers. We looked at the distribution of TD-errors as a function of time for a number of games (see Figure 9 in the appendix), and found that it becomes close to a heavy-tailed distribution as learning progresses, while still differing substantially across games; this empirically validates the form of Equation 1. Figure 10, in the appendix, shows how this distribution interacts with Equation 1 to produce the effective replay probabilities.

While doing this analysis, we stumbled upon another phenomenon (obvious in retrospect), namely that some fraction of the visited transitions are never replayed before they drop out of the sliding

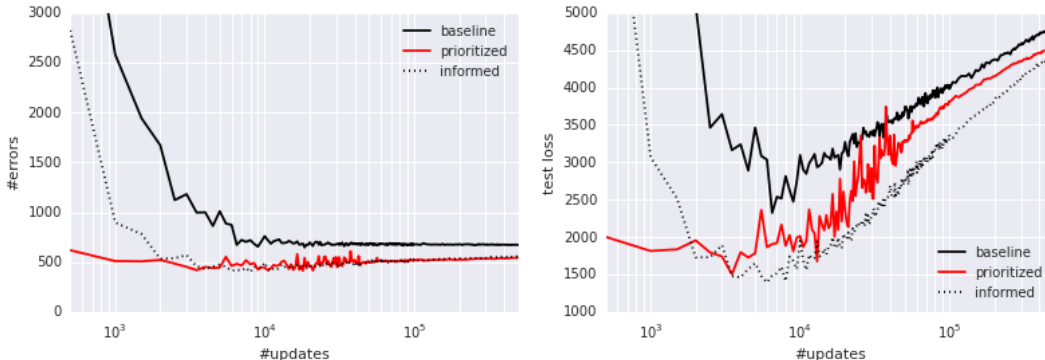


Figure 4: Classification errors as a function of supervised learning updates on severely class-imbalanced MNIST. Prioritized sampling improves performance, leading to comparable errors on the test set, and improving the imbalance-informed performance (median of 3 random initializations). **Left:** Number of misclassified test set samples. **Right:** Test set loss, highlighting overfitting.

window memory, and many more are first replayed only a long time after they are encountered. Also, uniform sampling is implicitly biased toward out-of-date transitions that were generated by a policy that has typically seen hundreds of thousands of updates since. Prioritized replay with its bonus for unseen transitions directly corrects the first of these issues, and also tends to help with the second one, as more recent transitions tend to have larger error – this is because old transitions will have had more opportunities to have them corrected, and because novel data tends to be less well predicted by the value function.

We hypothesize that deep neural networks interact with prioritized replay in another interesting way. When we distinguish learning the value given a representation (i.e., the top layers) from learning an improved representation (i.e., the bottom layers), then transitions for which the representation is good will quickly reduce their error and then be replayed much less, increasing the learning focus on others where the representation is poor, thus putting more resources into distinguishing aliased states – if the observations and network capacity allow for it.

## 6 EXTENSIONS

**Prioritized Supervised Learning:** The analogous approach to prioritized replay in the context of supervised learning is to sample non-uniformly from the dataset, each sample using a priority based on its last-seen error. This can help focus the learning on those samples that can still be learned from, devoting additional resources to the (hard) boundary cases, somewhat similarly to boosting (Galar et al., 2012). Furthermore, if the dataset is imbalanced, we hypothesize that samples from the rare classes will be sampled disproportionately often, because their errors shrink less fast, and the chosen samples from the common classes will be those nearest to the decision boundaries, leading to an effect similar to hard negative mining (Felzenszwalb et al., 2008). To check whether these intuitions hold, we conducted a preliminary experiment on a class-imbalanced variant of the classical MNIST digit classification problem (LeCun et al., 1998), where we removed 99% of the samples for digits 0, 1, 2, 3, 4 in the training set, while leaving the test/validation sets untouched (i.e., those retain class balance). We compare two scenarios: in the informed case, we reweight the errors of the impoverished classes artificially (by a factor 100), in the uninformed scenario, we provide no hint that the test distribution will be a different one. See Appendix B.3 for the details of the convolutional neural network training setup. Prioritized sampling (uninformed, with  $\alpha = 1, \beta = 0$ ) outperforms the uninformed uniform baseline, and approaches the performance of the informed uniform baseline in terms of generalization (see Figure 4); prioritized training is also faster in terms of learning speed.

**Off-policy Replay:** Two standard approaches to off-policy RL are rejection sampling, and using importance sampling ratios  $\rho$  to correct for how likely a transition would have been on-policy. Our approach contains analogues to both these approaches, the replay probability  $P$  and the IS-correction  $w$ . It appears therefore natural to apply it to off-policy RL, if transitions are available in a replay memory. In particular, we recover weighted IS with  $w = \rho, \alpha = 0, \beta = 1$  and rejection sampling with  $p = \min(1; \rho), \alpha = 1, \beta = 0$ , in the proportional variant. Our results indicate that intermediate



variants, possibly with annealing or ranking, could be more useful in practice – especially when IS ratios introduce high variance, i.e., when the policy of interest differs substantially from the behavior policy in some states. Of course, the off-policy correction is complementary to our prioritization based on expected learning progress, and the same framework can be used for a hybrid prioritization by defining  $p = \rho \cdot |\delta|$ , or some other sensible trade-off based on both  $\rho$  and  $\delta$ .

**Feedback for Exploration:** An interesting side-effect of prioritized replay is that the total number  $M_i$  that a transition will end up being replayed varies widely, and this gives a rough indication of how useful it was to the agent. This potentially valuable signal can be fed back to the exploration strategy that generates the transitions. For example, we could sample exploration hyper-parameters (such as the fraction of random actions  $\epsilon$ , the Boltzmann temperature, or the amount of intrinsic reward to mix in) from a parametrized distribution at the beginning of each episode, monitor the usefulness of the experience via  $M_i$ , and update the distribution toward generating more useful experience. Or, in a parallel system like the Gorilla agent (Nair et al., 2015), it could guide resource allocation between a collection of concurrent but heterogeneous ‘actors’, each with different exploration hyper-parameters.

**Prioritized Memories:** Considerations that help determine which transitions to replay are likely to be relevant for determining which memories to store and when to erase them (i.e., when it becomes unlikely that we would ever want to replay them anymore). An explicit control over which memories to keep or erase can help reduce the required total memory size, because it reduces redundancy (frequently visited transitions will have low error, so many of them will be dropped), while automatically adjusting for what has been learned already (dropping many of the ‘easy’ transitions) and biasing the contents of the memory to where the errors remain high. This is a non-trivial aspect, because memory requirements for DQN are currently dominated by the size of the replay memory, no longer by the size of the neural network. Erasing is a more final decision than reducing the replay probability, thus an even stronger emphasis of diversity may be necessary, for example by tracking the age of each transitions and using it to modulate the priority in such a way as to preserve sufficient old experience to prevent cycles (related to ‘hall of fame’ ideas in multi-agent literature, Rosin & Belew, 1997) or collapsing value functions. The priority mechanism is also flexible enough to permit integrating experience from *other sources*, such as from a planner or from human expert trajectories (Guo et al., 2014), since knowing the source can be used to modulate each transition’s priority, for example in such a way as to preserve a sufficient fraction of external experience in memory.

## 7 CONCLUSION

This paper introduced prioritized replay, a method that can make learning from experience replay more efficient. We studied a couple of variants, devised implementations that scale to large replay memories, and found that prioritized replay speeds up learning and leads to a new state-of-the-art on the Atari benchmark. We laid out further variants and extensions that hold promise, namely for class-imbalanced supervised learning.

## ACKNOWLEDGMENTS

We thank our Deepmind colleagues, in particular Hado van Hasselt, Joseph Modayil, Nicolas Heess, Marc Bellemare, Razvan Pascanu, Dharshan Kumaran, Daan Wierstra, Arthur Guez, Charles Blundell, Alex Pritzel, Alex Graves, Balaji Lakshminarayanan, Ziyu Wang, Nando de Freitas, Remi Munos and Geoff Hinton for insightful discussions and feedback.

## REFERENCES

- Andre, David, Friedman, Nir, and Parr, Ronald. Generalized prioritized sweeping. In *Advances in Neural Information Processing Systems*. Citeseer, 1998.
- Atherton, Laura A, Dupret, David, and Mellor, Jack R. Memory trace replay: the shaping of memory consolidation by neuromodulation. *Trends in neurosciences*, 38(9):560–570, 2015.
- Bellemare, Marc G, Naddaf, Yavar, Veness, Joel, and Bowling, Michael. The arcade learning environment: An evaluation platform for general agents. *arXiv preprint arXiv:1207.4708*, 2012.

- Bellemare, Marc G., Ostrovski, Georg, Guez, Arthur, Thomas, Philip S., and Munos, Rémi. Increasing the action gap: New operators for reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2016.
- Collobert, Ronan, Kavukcuoglu, Koray, and Farabet, Clément. Torch7: A matlab-like environment for machine learning. In *BigLearn, NIPS Workshop*, number EPFL-CONF-192376, 2011.
- Diamond, Jared. Zebras and the Anna Karenina principle. *Natural History*, 103:4–4, 1994.
- Felzenszwalb, Pedro, McAllester, David, and Ramanan, Deva. A discriminatively trained, multi-scale, deformable part model. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pp. 1–8. IEEE, 2008.
- Foster, David J and Wilson, Matthew A. Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, 440(7084):680–683, 2006.
- Galar, Mikel, Fernandez, Alberto, Barrenechea, Edurne, Bustince, Humberto, and Herrera, Francisco. A review on ensembles for the class imbalance problem: bagging-, boosting-, and hybrid-based approaches. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 42(4):463–484, 2012.
- Geramifard, Alborz, Doshi, Finale, Redding, Joshua, Roy, Nicholas, and How, Jonathan. Online discovery of feature dependencies. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pp. 881–888, 2011.
- Guo, Xiaoxiao, Singh, Satinder, Lee, Honglak, Lewis, Richard L, and Wang, Xiaoshi. Deep Learning for Real-Time Atari Game Play Using Offline Monte-Carlo Tree Search Planning. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., and Weinberger, K.Q. (eds.), *Advances in Neural Information Processing Systems 27*, pp. 3338–3346. Curran Associates, Inc., 2014.
- Hinton, Geoffrey E. To recognize shapes, first learn to generate images. *Progress in brain research*, 165:535–547, 2007.
- Kingma, Diederik P. and Ba, Jimmy. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.
- Lecun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, Nov 1998. ISSN 0018-9219. doi: 10.1109/5.726791.
- LeCun, Yann, Cortes, Corinna, and Burges, Christopher JC. The MNIST database of handwritten digits, 1998.
- Lin, Long-Ji. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine learning*, 8(3-4):293–321, 1992.
- Mahmood, A Rupam, van Hasselt, Hado P, and Sutton, Richard S. Weighted importance sampling for off-policy learning with linear function approximation. In *Advances in Neural Information Processing Systems*, pp. 3014–3022, 2014.
- McNamara, Colin G, Tejero-Cantero, Álvaro, Trouche, Stéphanie, Campo-Urriza, Natalia, and Dupret, David. Dopaminergic neurons promote hippocampal reactivation and spatial memory persistence. *Nature neuroscience*, 2014.
- Mnih, Volodymyr, Kavukcuoglu, Koray, Silver, David, Graves, Alex, Antonoglou, Ioannis, Wierstra, Daan, and Riedmiller, Martin. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Mnih, Volodymyr, Kavukcuoglu, Koray, Silver, David, Rusu, Andrei A, Veness, Joel, Bellemare, Marc G, Graves, Alex, Riedmiller, Martin, Fidjeland, Andreas K, Ostrovski, Georg, Petersen, Stig, Beattie, Charles, Sadik, Amir, Antonoglou, Ioannis, King, Helen, Kumaran, Dharshan, Wierstra, Daan, Legg, Shane, and Hassabis, Demis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- Moore, Andrew W and Atkeson, Christopher G. Prioritized sweeping: Reinforcement learning with less data and less time. *Machine Learning*, 13(1):103–130, 1993.
- Nair, Arun, Srinivasan, Praveen, Blackwell, Sam, Alcicek, Cagdas, Fearon, Rory, Maria, Alessandro De, Panneershelvam, Vedavyas, Suleyman, Mustafa, Beattie, Charles, Petersen, Stig, Legg, Shane, Mnih, Volodymyr, Kavukcuoglu, Koray, and Silver, David. Massively parallel methods for deep reinforcement learning. *arXiv preprint arXiv:1507.04296*, 2015.
- Ólafsdóttir, H Freyja, Barry, Caswell, Saleem, Aman B, Hassabis, Demis, and Spiers, Hugo J. Hippocampal place cells construct reward related sequences through unexplored space. *Elife*, 4:

- e06063, 2015.
- Riedmiller, Martin. Rprop-description and implementation details. 1994.
- Rosin, Christopher D and Belew, Richard K. New methods for competitive coevolution. *Evolutionary Computation*, 5(1):1–29, 1997.
- Schaul, Tom, Zhang, Sixin, and Lecun, Yann. No more pesky learning rates. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pp. 343–351, 2013.
- Schmidhuber, Jürgen. Curious model-building control systems. In *Neural Networks, 1991. 1991 IEEE International Joint Conference on*, pp. 1458–1463. IEEE, 1991.
- Singer, Annabelle C and Frank, Loren M. Rewarded outcomes enhance reactivation of experience in the hippocampus. *Neuron*, 64(6):910–921, 2009.
- Stadie, Bradley C, Levine, Sergey, and Abbeel, Pieter. Incentivizing exploration in reinforcement learning with deep predictive models. *arXiv preprint arXiv:1507.00814*, 2015.
- Sun, Yi, Ring, Mark, Schmidhuber, Jürgen, and Gomez, Faustino J. Incremental basis construction from temporal difference error. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pp. 481–488, 2011.
- van Hasselt, Hado. Double Q-learning. In *Advances in Neural Information Processing Systems*, pp. 2613–2621, 2010.
- van Hasselt, Hado, Guez, Arthur, and Silver, David. Deep Reinforcement Learning with Double Q-learning. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, 2016. URL <http://arxiv.org/abs/1509.06461>.
- van Seijen, Harm and Sutton, Richard. Planning by prioritized sweeping with small backups. In *Proceedings of The 30th International Conference on Machine Learning*, pp. 361–369, 2013.
- Wang, Z., de Freitas, N., and Lanctot, M. Dueling network architectures for deep reinforcement learning. Technical report, 2015.
- Watkins, Christopher JCH and Dayan, Peter. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.
- White, Adam, Modayil, Joseph, and Sutton, Richard S. Surprise and curiosity for big data robotics. In *Workshops at the Twenty-Eighth AAAI Conference on Artificial Intelligence*, 2014.

## A PRIORITIZATION VARIANTS

The absolute TD-error is only one possible proxy for the ideal priority measure of expected learning progress. While it captures the scale of potential improvement, it ignores inherent stochasticity in rewards or transitions, as well as possible limitations from partial observability or FA capacity; in other words, it is problematic when there are unlearnable transitions. In that case, its *derivative* – which could be approximated by the difference between a transition’s current  $|\delta|$  and the  $|\delta|$  when it was last replayed<sup>3</sup> – may be more useful. This measure is less immediately available, however, and is influenced by whatever was replayed in the meanwhile, which increases its variance. In preliminary experiments, we found it did not outperform  $|\delta|$ , but this may say more about the class of (near-deterministic) environments we investigated, than about the measure itself.

An orthogonal variant is to consider the norm of the *weight-change* induced by replaying a transition – this can be effective if the underlying optimizer employs adaptive step-sizes that reduce gradients in high-noise directions (Schaul et al., 2013; Kingma & Ba, 2014), thus placing the burden of distinguishing between learnable and unlearnable transitions on the optimizer.

It is possible to modulate prioritization by not treating positive TD-errors the same than negative ones; we can for example invoke the Anna Karenina principle (Diamond, 1994), interpreted to mean that there are many ways in which a transition can be less good than expected, but only one in which can be better, to introduce an *asymmetry* and prioritize positive TD-errors more than negative ones of equal magnitude, because the former are more likely to be informative. Such an asymmetry in replay frequency was also observed in rat studies (Singer & Frank, 2009). Again, our preliminary experiments with such variants were inconclusive.

The evidence from neuroscience suggest that a prioritization based on episodic return rather than expected learning progress may be useful too Atherton et al. (2015); Ólafsdóttir et al. (2015); Foster & Wilson (2006). For this case, we could boost the replay probabilities of entire episodes, instead of transitions, or boost individual transitions by their observed return-to-go (or even their value estimates).

For the issue of preserving sufficient diversity (to prevent overfitting, premature convergence or impoverished representations), there are alternative solutions to our choice of introducing stochasticity, for example, the priorities could be modulated by a novelty measure in observation space. Nothing prevents a hybrid approach where some fraction of the elements (of each minibatch) are sampled according to one priority measure and the rest according to another one, introducing additional diversity. An orthogonal idea is to increase priorities of transitions that have not been replayed for a while, by introducing an explicit *staleness bonus* that guarantees that every transition is revisited from time to time, with that chance increasing at the same rate as its last-seen TD-error becomes stale. In the simple case where this bonus grows linearly with time, this can be implemented at no additional cost by subtracting a quantity proportional to the global step-count from the new priority on any update.<sup>4</sup>

In the particular case of RL with bootstrapping from value functions, it is possible to exploit the sequential structure of the replay memory using the following intuition: a transition that led to a large amount of learning (about its outgoing state) has the potential to change the bootstrapping target for all transitions leading into that state, and thus there is more to be learned about these. Of course we know at least one of these, namely the historic *predecessor* transition, and so boosting its priority makes it more likely to be revisited soon. Similarly to eligibility traces, this lets information trickle backward from a future outcome to the value estimates of the actions and states that led to it. In practice, we add  $|\delta|$  of the current transition to predecessor transition’s priority, but only if the predecessor transition is not a terminal one. This idea is related to ‘reverse replay’ observed in rodents Foster & Wilson (2006), and to a recent extension of prioritized sweeping (van Seijen & Sutton, 2013).

<sup>3</sup>Of course, more robust approximations would be a function of the history of all encountered  $\delta$  values. In particular, one could imagine an RProp-style update (Riedmiller, 1994) to priorities that increase the priority while the signs match, and reduce it whenever consecutive errors (for the same transition) differ in sign.

<sup>4</sup>If bootstrapping is used with policy iteration, such that the target values come from separate network (as is the case for DQN), then there is a large increase in uncertainty about the priorities when the target network is updated in the outer iteration. At these points, the staleness bonus is increased in proportion to the number of individual (low-level) updates that happened in-between.

## B EXPERIMENTAL DETAILS

### B.1 BLIND CLIFFWALK

For the Blind Cliffwalk experiments (Section 3.1 and following), we use a straight-forward Q-learning (Watkins & Dayan, 1992) setup. The Q-values are represented using either a tabular look-up table, or a linear function approximator, in both cases represented  $Q(s, a) := \theta^\top \phi(s, a)$ . For each transition, we compute its TD-error using:

$$\delta_t := R_t + \gamma_t \max_a Q(S_t, a) - Q(S_{t-1}, A_{t-1}) \quad (2)$$

and update the parameters using stochastic gradient ascent:

$$\theta \leftarrow \theta + \eta \cdot \delta_t \cdot \nabla_{\theta} Q|_{S_{t-1}, A_{t-1}} = \theta + \eta \cdot \delta_t \cdot \phi(S_{t-1}, A_{t-1}) \quad (3)$$

For the linear FA case we use a very simple encoding of state as a 1-hot vector (as for tabular), but concatenated with a constant bias feature of value 1. To make generalization across actions impossible, we alternate which action is ‘right’ and which one is ‘wrong’ for each state. All elements are initialized with small values near zero,  $\theta_i \sim \mathcal{N}(0, 0.1)$ .

We vary the size of the problem (number of states  $n$ ) from 2 to 16. The discount factor is set to  $\gamma = 1 - \frac{1}{n}$  which keeps values on approximately the same scale independently of  $n$ . This allows us to use a fixed step-size of  $\eta = \frac{1}{4}$  in all experiments.

The replay memory is filled by exhaustively executing all  $2^n$  possible sequences of actions until termination (in random order). This guarantees that exactly one sequence will succeed and hit the final reward, and all others will fail with zero reward. The replay memory contains all the relevant experience (the total number of transitions is  $2^{n+1} - 2$ ), at the frequency that it would be encountered when acting online with a random behavior policy. Given this, we can in principle learn until convergence by just increasing the amount of computation; here, convergence is defined as a mean-squared error (MSE) between the Q-value estimates and the ground truth below  $10^{-3}$ .

### B.2 ATARI EXPERIMENTS

#### B.2.1 IMPLEMENTATION DETAILS

Prioritizing using a replay memory with  $N = 10^6$  transitions introduced some performance challenges. Here we describe a number of things we did to minimize additional run-time and memory overhead, extending the discussion in Section 3.

**Rank-based prioritization** Early experiments with Atari showed that maintaining a sorted data structure of  $10^6$  transitions with constantly changing TD-errors dominated running time. Our final solution was to store transitions in a priority queue implemented using an array-based binary heap. The underlying heap array was directly used as an approximation of a sorted array. This is unconventional, however our tests on smaller environments showed learning was unaffected compared to using a perfectly sorted array. Likely this is in part due to TD-errors already being an proxy for the usefulness of a transition and our use of stochastic prioritized sampling. Another small improvement came from avoiding recalculation of partitions for the sampling distribution. We reused the same partition for values of  $N$  that are close together and by updating  $\alpha$  and  $\beta$  infrequently where possible. Our final implementation for rank-based prioritization produced an additional 2%-4% increase in running time and negligible additional memory usage. This could be reduced further in a number of ways, e.g. with a more efficient heap implementation, but was good enough for our experiments.

**Proportional prioritization** The ‘sum-tree’ data structure used here is very similar in spirit to the array representation of a binary heap. However, instead of the usual heap property, the value of a parent node is the sum of its children. Leaf nodes store the transition priorities and the internal nodes are intermediate sums, with the parent node containing the sum over all priorities,  $p_{\text{total}}$ . This provides a efficient way of calculating the cumulative sum of priorities, allowing  $O(\log N)$  updates and sampling. To sample a minibatch of size  $k$ , the range  $[0, p_{\text{total}}]$  is divided equally into  $k$  ranges. Next, a value is uniformly sampled from each range. Finally the transitions that correspond to each of these sampled values are retrieved from the tree. Overhead is similar to rank-based prioritization.

As mentioned in Section 3.4, whenever importance sampling is used, all weights  $w_i$  were scaled so that  $\max_i w_i = 1$ . We found that this worked better in practice as it kept all weights within a reasonable range, avoiding the possibility of extremely large updates. It is worth mentioning that this normalization interacts with annealing on  $\beta$ : as  $\beta$  approaches 1, the normalization constant grows, which reduces the effective average update in a similar way to annealing the step-size  $\eta$ .

### B.2.2 HYPER-PARAMETERS

Throughout this paper our baseline was DQN and the tuned version of Double DQN. Tuning of hyper-parameters was done over a subset of Atari games (Breakout, Pong, Ms. Pac-Man, Q\*bert, Alien, Battlezone, Asterix). Table 2 lists the values tried and Table 3 lists the chosen parameters.

Hyper-parameter	Range of values
$\alpha$	0, 0.4, 0.5, 0.6, 0.7, 0.8
$\beta$	0, 0.4, 0.5, 0.6, 1
$\eta$	$\eta_{\text{baseline}}, \eta_{\text{baseline}}/2, \eta_{\text{baseline}}/4, \eta_{\text{baseline}}/8$

Table 2: Hyper-parameters considered in experiments. Here  $\eta_{\text{baseline}} = 0.00025$ .

Hyper-parameter	DQN		Double DQN		
	baseline	rank-based	baseline	rank-based	proportional
$\alpha$ (priority)	0	0.5 $\rightarrow$ 0	0	0.7	0.6
$\beta$ (IS)	0	0	0	0.5 $\rightarrow$ 1	0.4 $\rightarrow$ 1
$\eta$ (step-size)	0.00025	$\eta_{\text{baseline}}/4$	0.00025	$\eta_{\text{baseline}}/4$	$\eta_{\text{baseline}}/4$

Table 3: Chosen hyper-parameters for prioritized variants of DQN. Arrows indicate linear annealing, where the limiting value is reached at the end of training. Note the rank-based variant with DQN as the baseline is an early version without IS. Here, the bias introduced by prioritized replay was instead corrected by annealing  $\alpha$  to zero.

### B.2.3 EVALUATION

We evaluate our agents using *human starts* and *no-ops* evaluation methods as described in (van Hasselt et al., 2016). Human starts evaluation uses start states sampled randomly from human traces. No-ops evaluation uses start states that are randomized by doing a random number of no-ops at the beginning of each episode. This is close to the *test* evaluation that the agents periodically use during training. For both human starts and no-ops, the score is averaged over 100 evaluations of 30 minutes each. All learning curves plots show the scores generated using the test evaluation method.

Table 4 and Table 5 show evaluation method differences and the  $\epsilon$  used in the  $\epsilon$ -greedy policy for each agent during evaluation. The agent used for the no-ops and human starts evaluation is the best agent seen during training as was done in (van Hasselt et al., 2016).

Evaluation method	Frames	Emulator time	Number of evaluations	Agent start point
Human starts	108,000	30 mins	100	human starts
No-ops	108,000	30 mins	100	up to 30 random no-ops
Test	500,000	139 mins	1	up to 30 random no-ops

Table 4: Evaluation method comparison.

Evaluation method	DQN		Double DQN		
	baseline	rank-based	baseline	rank-based	proportional
Human starts	0.05	0.01	0.001	0.001	0.001
No-ops	0.05	0.01	0.01	0.001	0.001
Test	0.05	0.05	0.01	0.01	0.01

Table 5: The  $\epsilon$  used in the  $\epsilon$ -greedy policy for each agent, for evaluation method.

Normalized score is calculated as in (van Hasselt et al., 2016):

$$\text{SCORE}_{\text{normalized}} = \frac{\text{SCORE}_{\text{agent}} - \text{SCORE}_{\text{random}}}{|\text{SCORE}_{\text{human}} - \text{SCORE}_{\text{random}}|} \quad (4)$$

Note the absolute value of the denominator is taken. This only affects Video Pinball where the random score is higher than the human score. Combined with a high agent score, Video Pinball has a very large effect on the mean normalized score. We continue to use this so our normalized scores are comparable.

### B.3 CLASS-IMBALANCED MNIST

#### B.3.1 DATASET SETUP

In our supervised learning setting we modified MNIST to obtain a new training dataset with a significant label imbalance. This new dataset was obtained by considering a small subset of the samples that correspond to the first 5 digits (0, 1, 2, 3, 4) and all of the samples that correspond to the remaining 5 labels (5, 6, 7, 8, 9). For each of the first 5 digits we randomly sampled 1% of the available examples, i.e., 1% of the available 0s, 1% of the available 1s etc. In the resulting dataset there are examples of all 10 different classes but it is highly imbalanced since there are 100 times more examples that correspond to the 5, 6, 7, 8, 9 classes than to the 0, 1, 2, 3, 4 ones. In all our experiments we used the original MNIST *test* dataset without removing any samples.

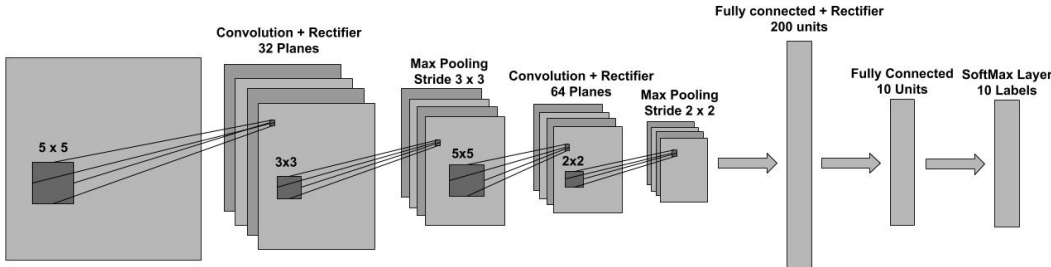


Figure 5:

The architecture of the feed-forward network used in the Prioritized Supervised Learning experiments.

#### B.3.2 TRAINING SETUP

In our experiments we used a 4 layer feed-forward neural network with an architecture similar to that of LeNet5 (Lecun et al., 1998). This is a 2 layer convolutional neural network followed by 2 fully connected layers at the top. Each convolutional layer is comprised of a pure convolution followed by a rectifier non-linearity and a subsampling max pooling operation. The two fully-connected layers in the network are also separated by a rectifier non-linearity. The last layer is a softmax which is used to obtain a normalized distribution over the possible labels. The complete architecture is shown in Figure 5, and is implemented using Torch7 (Collobert et al., 2011). The model was trained using stochastic gradient descent with no momentum and a minibatch of size 60. In all our experiments we considered 6 different step-sizes (0.3, 0.1, 0.03, 0.01, 0.003 and 0.001) and for each case presented in this work, we selected the step-size that lead to the best (balanced) validation performance. We used the negative log-likelihood loss criterion and we ran experiments with both the weighted and unweighted version of the loss. In the weighted case the loss of the examples that correspond to the first 5 digits (0, 1, 2, 3, 4) was scaled by a factor of a 100 to accommodate the label imbalance in the training set described above.

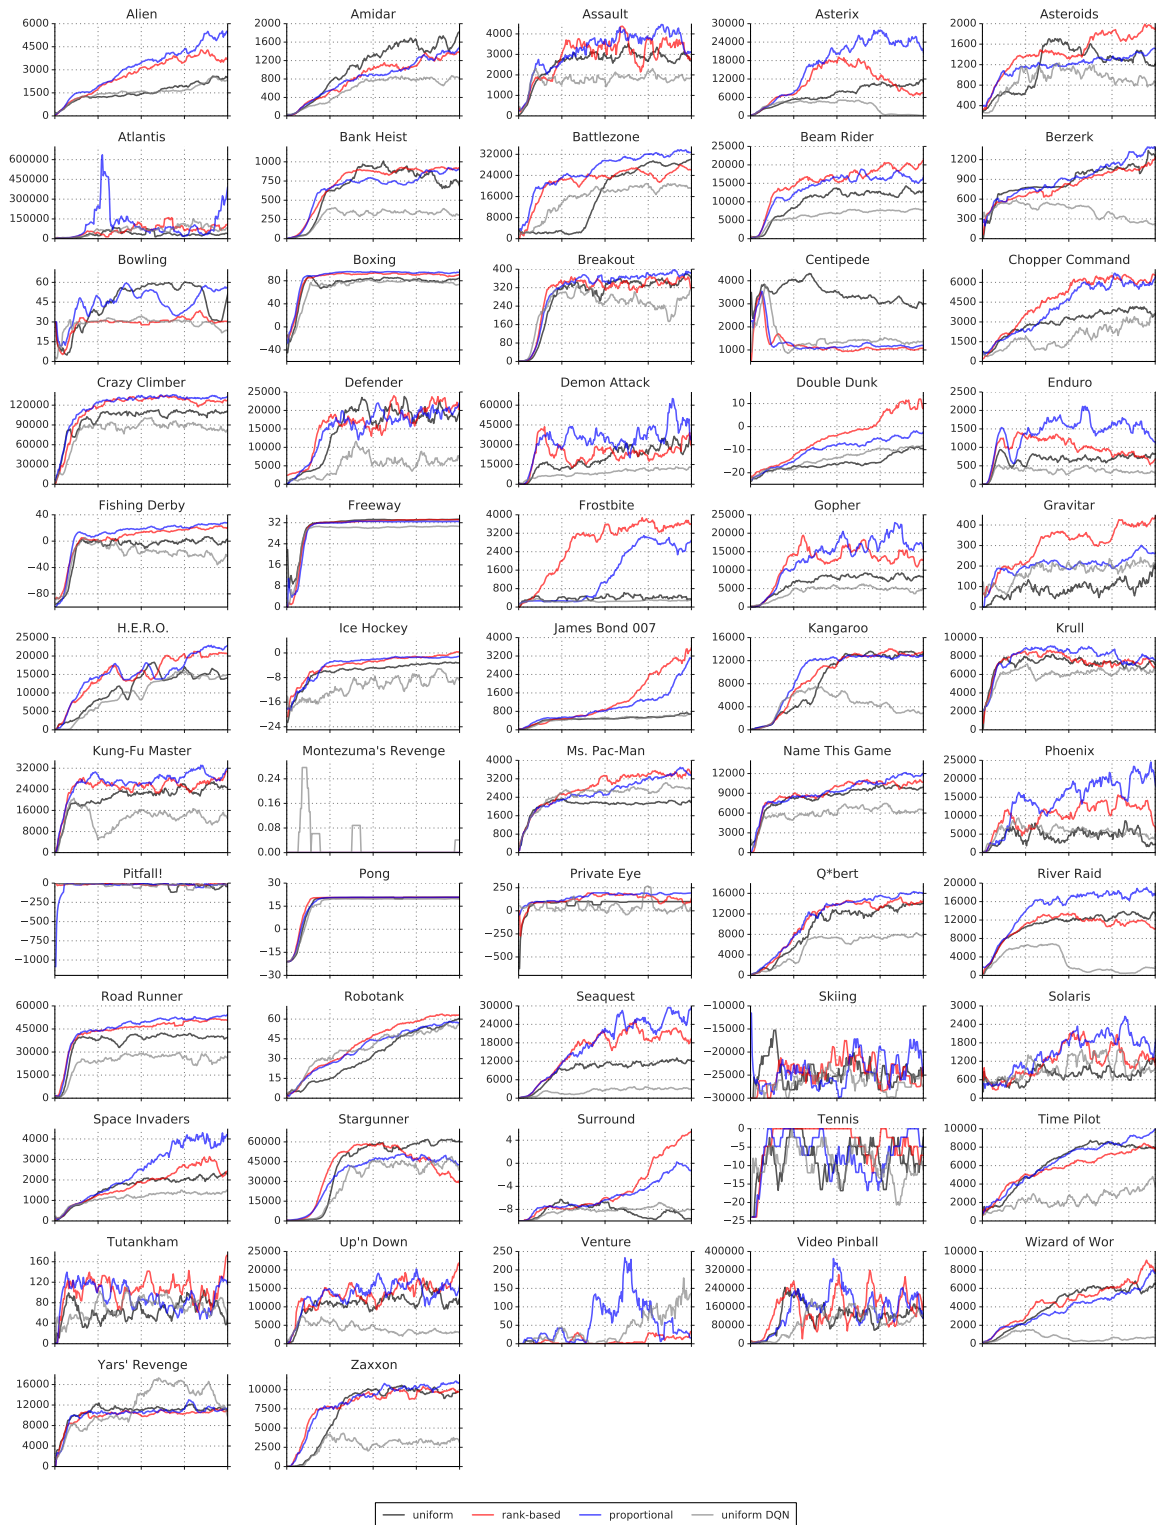


Figure 6: Learning curves (in raw score) for Double DQN (uniform baseline, in black), with rank-based prioritized replay (red), proportional prioritization (blue), for all 57 games of the Atari benchmark suite. Each curve corresponds to a single training run over 200 million frames, with the test evaluation method described in Section B.2.3, smoothed with a moving average with window size 10. In gray are the learning curves for the original DQN. See also Figure 7 for a more detailed view of a selection of these.



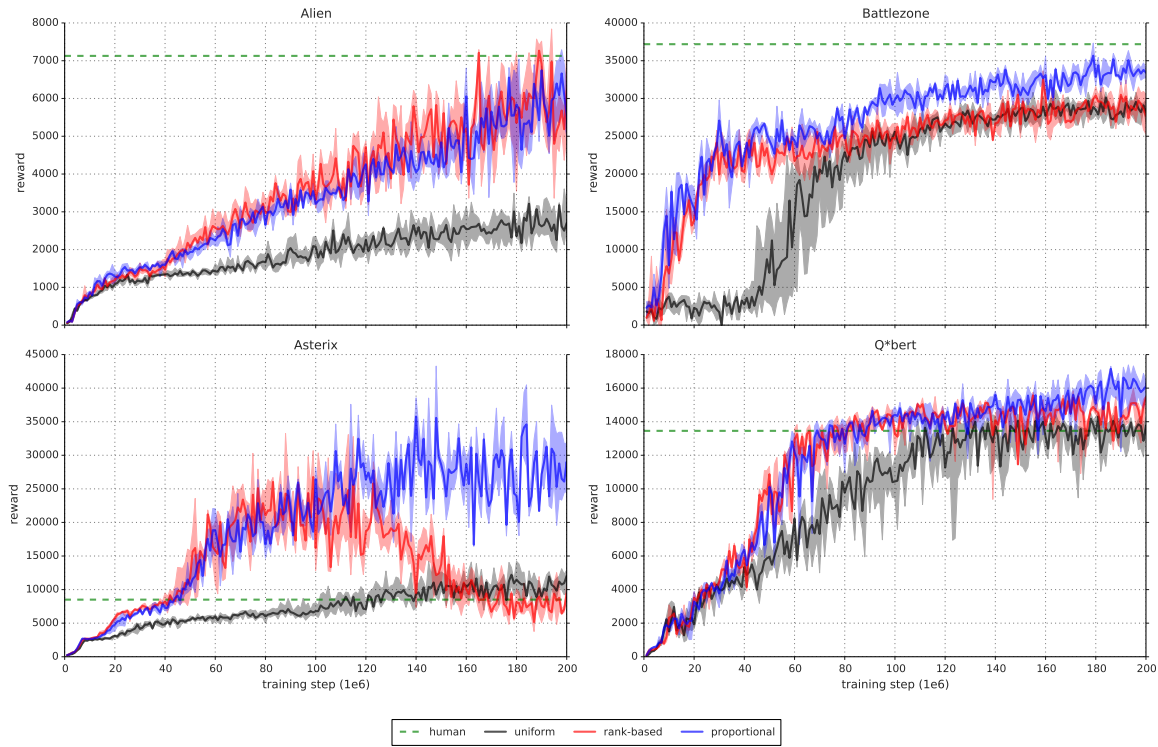


Figure 7: Detailed learning curves for rank-based (red) and proportional (blue) prioritization, as compared to the uniform Double DQN baseline (black) on a selection of games. The solid lines are the median scores, and the shaded area denotes the interquartile range across 4 or more random initializations. While the variability between runs is substantial, there are significant differences in final achieved score, and also in learning speed.

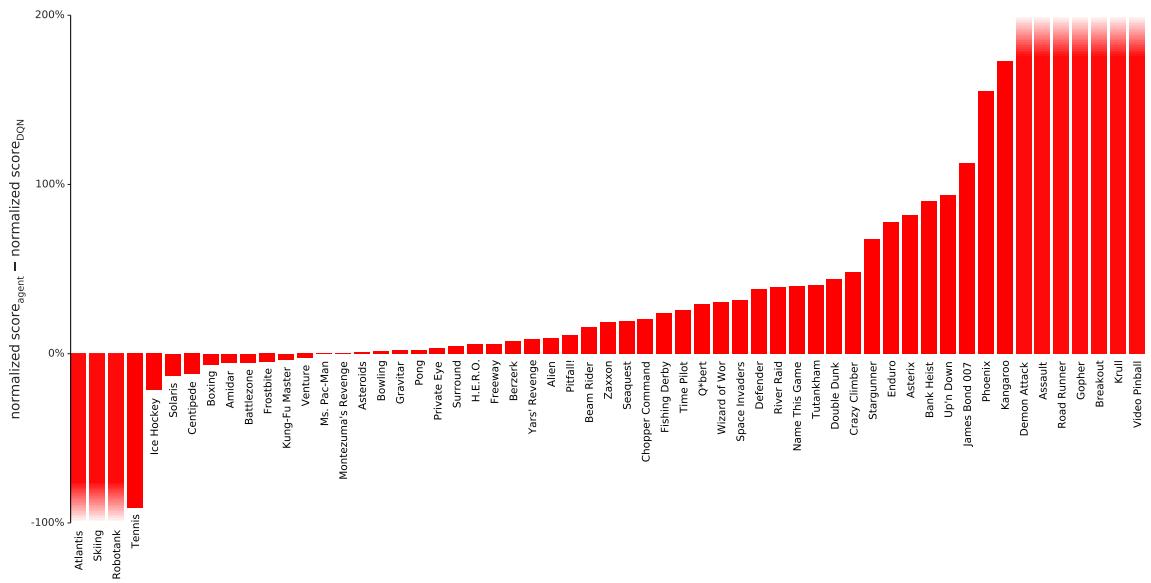


Figure 8: Difference in normalized score (the gap between random and human is 100%) on 57 games with human starts, comparing DQN with and without rank-based prioritized replay, showing substantial improvements in many games. Exact scores are in Table 6. See also Figure 3 where Double DQN is the baseline.

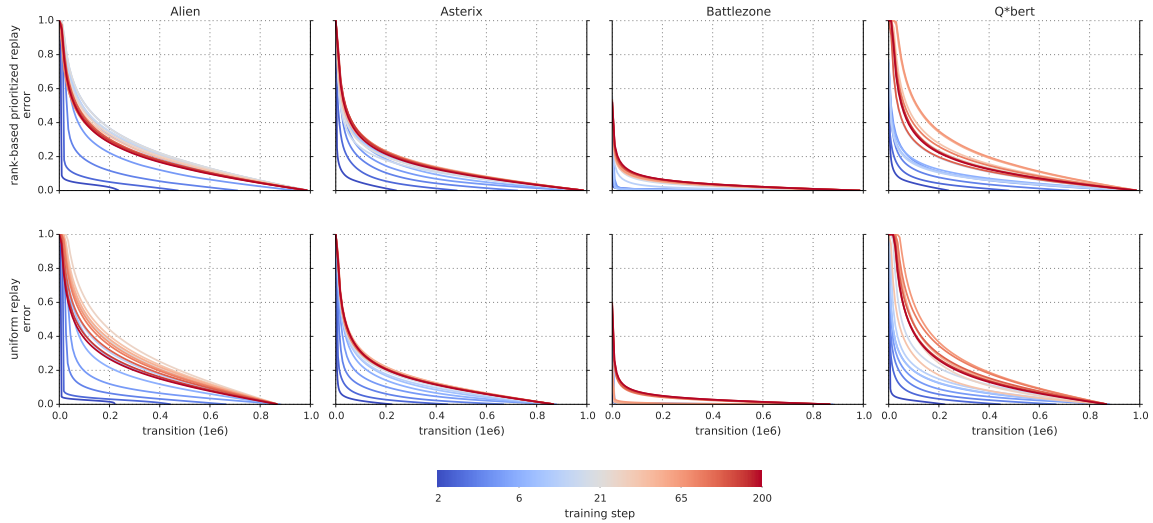


Figure 9: Visualization of the last-seen absolute TD-error for all transitions in the replay memory, sorted, for a selection of Atari games. The lines are color-coded by the time during learning, at a resolution of  $10^6$  frames, with the coldest colors in the beginning and the warmest toward the end of training. We observe that in some games it starts quite peaked but quickly becomes spread out, following approximately a heavy-tailed distribution. This phenomenon happens for both rank-based prioritized replay (top) and uniform replay (bottom) but is faster for prioritized replay.

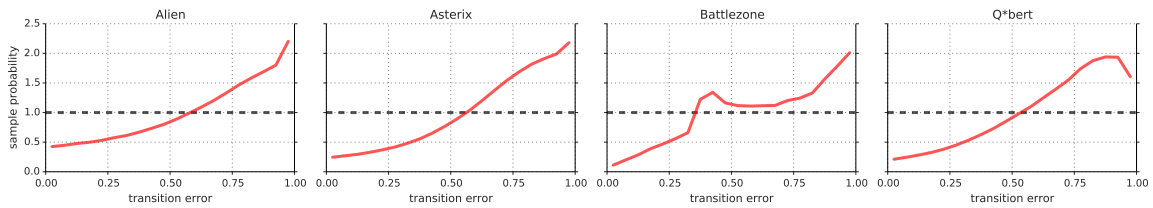


Figure 10: **Effective replay probability**, as a function of absolute TD-error, for the rank-based prioritized replay variant near the start of training. This shows the effect of Equation 1 with  $\alpha = 0.7$  in practice, compared to the uniform baseline (dashed horizontal line). The effect is irregular, but qualitatively similar for a selection of games.

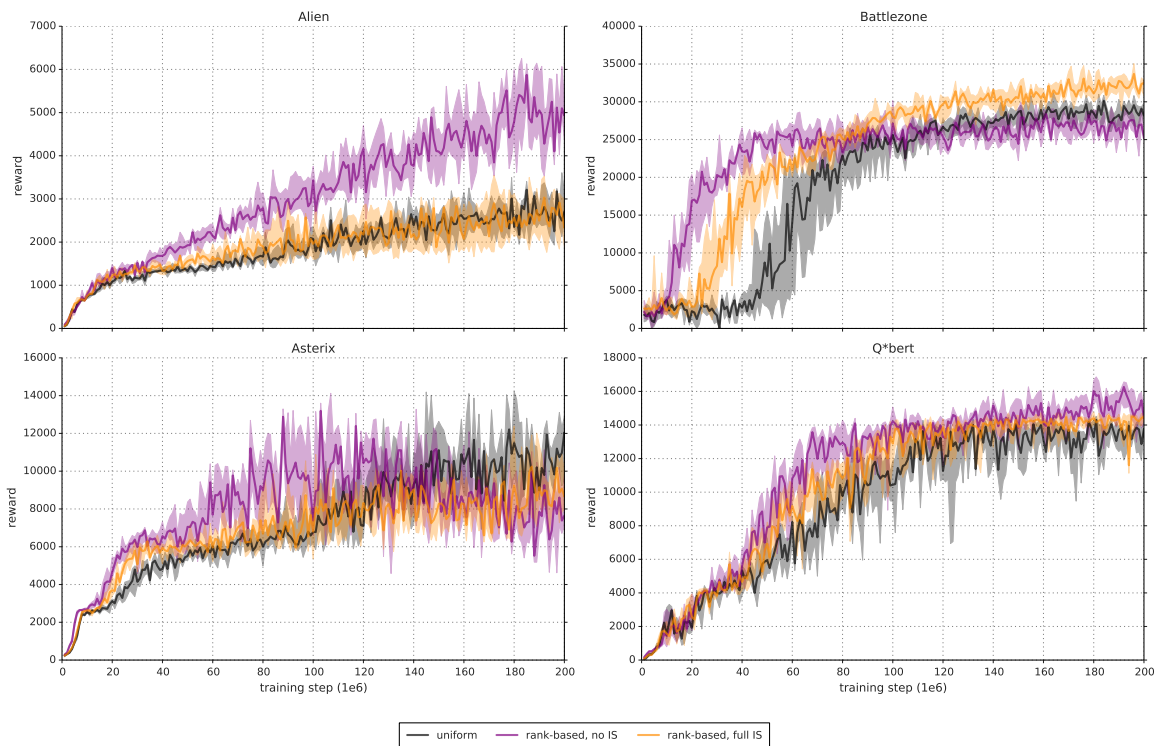


Figure 11: **Effect of importance sampling:** These learning curves (as in Figure 7) show how rank-based prioritization is affected by full importance-sampling correction (i.e.,  $\beta = 1$ , in orange), as compared to the uniform baseline (black,  $\alpha = 0$ ) and pure, uncorrected prioritized replay (violet,  $\beta = 0$ ), on a few selected games. The shaded area corresponds to the interquartile range. The step-size for full IS correction is the same as for uniform replay. For uncorrected prioritized replay, the step-size is reduced by a factor of 4. Compared to uncorrected prioritized replay, importance sampling makes learning less aggressive, leading on the one hand to slower initial learning, but on the other hand to a smaller risk of premature convergence and sometimes better ultimate results. Compared to uniform replay, fully corrected prioritization is on average better.

Game	DQN		Double DQN		
	baseline	rank-based	baseline	rank-based	proportional
Alien	8%	17%	14%	<b>19%</b>	12%
Amidar	11%	6%	10%	8%	<b>14%</b>
Assault	326%	631%	1248%	1381%	<b>1641%</b>
Asterix	41%	123%	226%	303%	<b>431%</b>
Asteroids	2%	2%	1%	<b>2%</b>	2%
Atlantis	2049%	1480%	1481%	2419%	<b>4425%</b>
Bank Heist	39%	129%	117%	<b>137%</b>	128%
Battlezone	69%	63%	72%	75%	<b>87%</b>
Beam Rider	65%	80%	117%	<b>210%</b>	176%
Berzerk	15%	22%	39%	33%	<b>47%</b>
Bowling	19%	20%	<b>31%</b>	15%	27%
Boxing	647%	641%	<b>676%</b>	665%	632%
Breakout	1342%	<b>1823%</b>	1397%	1298%	1407%
Centipede	<b>24%</b>	12%	22%	19%	18%
Chopper Command	53%	<b>73%</b>	27%	48%	72%
Crazy Climber	381%	429%	392%	507%	<b>522%</b>
Defender	113%	151%	<b>207%</b>	176%	155%
Demon Attack	382%	596%	1698%	1888%	<b>2256%</b>
Double Dunk	625%	669%	781%	<b>2000%</b>	1169%
Enduro	86%	164%	154%	233%	<b>239%</b>
Fishing Derby	74%	98%	96%	<b>106%</b>	105%
Freeway	105%	111%	113%	<b>113%</b>	109%
Frostbite	10%	5%	33%	<b>83%</b>	69%
Gopher	373%	836%	728%	1679%	<b>2792%</b>
Gravitar	2%	<b>4%</b>	-2%	1%	-1%
H.E.R.O.	50%	56%	55%	<b>80%</b>	78%
Ice Hockey	79%	58%	71%	<b>93%</b>	85%
James Bond 007	198%	311%	161%	<b>1172%</b>	1038%
Kangaroo	167%	339%	416%	<b>458%</b>	384%
Krull	528%	<b>1051%</b>	590%	598%	653%
Kung-Fu Master	100%	97%	146%	<b>153%</b>	151%
Montezuma's Revenge	1%	0%	0%	<b>1%</b>	-0%
Ms. Pac-Man	5%	5%	5%	<b>11%</b>	11%
Name This Game	99%	138%	142%	173%	<b>200%</b>
Phoenix	114%	270%	202%	284%	<b>474%</b>
Pitfall!	-8%	2%	3%	-1%	<b>5%</b>
Pong	107%	110%	110%	<b>110%</b>	<b>110%</b>
Private Eye	-1%	<b>2%</b>	-2%	0%	-1%
Q*bert	76%	<b>106%</b>	87%	82%	93%
River Raid	30%	70%	74%	81%	<b>128%</b>
Road Runner	524%	<b>854%</b>	643%	780%	850%
Robotank	866%	752%	<b>872%</b>	828%	815%
Seaquest	10%	29%	36%	63%	<b>97%</b>
Skiing	27%	-122%	33%	<b>44%</b>	38%
Solaris	-8%	-21%	-14%	<b>3%</b>	2%
Space Invaders	87%	118%	191%	291%	<b>693%</b>
Stargunner	592%	660%	653%	<b>689%</b>	580%
Surround	25%	29%	75%	<b>103%</b>	58%
Tennis	<b>221%</b>	130%	93%	110%	132%
Time Pilot	64%	89%	140%	113%	<b>176%</b>
Tutankham	26%	<b>67%</b>	63%	35%	17%
Up'n Down	80%	173%	200%	125%	<b>313%</b>
Venture	12%	9%	0%	7%	<b>22%</b>
Video Pinball	2785%	4042%	7221%	5727%	<b>7367%</b>
Wizard of Wor	21%	52%	144%	131%	<b>177%</b>
Yars' Revenge	3%	<b>11%</b>	10%	7%	10%
Zaxxon	49%	68%	101%	113%	<b>113%</b>

Table 6: Normalized scores on 57 Atari games (random is 0%, human is 100%), from a single training run each, using human starts evaluation (see Section B.2.3). Baselines are from van Hasselt et al. (2016), see Equation 4 for how normalized scores are calculated.

Game	DQN		Double DQN		
	baseline	rank-based	baseline	rank-based	proportional
Alien	20%	67%	51%	58%	<b>93%</b>
Amidar	57%	47%	104%	107%	<b>119%</b>
Assault	781%	619%	995%	1434%	<b>1490%</b>
Asterix	50%	126%	207%	378%	<b>495%</b>
Asteroids	1%	2%	0%	4%	2%
Atlantis	1651%	1351%	576%	2129%	<b>2564%</b>
Bank Heist	60%	150%	138%	141%	<b>151%</b>
Battlezone	79%	62%	84%	84%	<b>103%</b>
Beam Rider	50%	60%	81%	<b>139%</b>	133%
Berzerk	18%	29%	44%	47%	<b>59%</b>
Bowling	20%	18%	<b>33%</b>	18%	29%
Boxing	732%	<b>828%</b>	762%	796%	822%
Breakout	1333%	<b>1462%</b>	1447%	1292%	1319%
Centipede	26%	13%	<b>33%</b>	24%	31%
Chopper Command	81%	89%	76%	<b>118%</b>	66%
Crazy Climber	399%	449%	425%	521%	<b>688%</b>
Defender	131%	141%	<b>205%</b>	180%	135%
Demon Attack	660%	1164%	3183%	<b>3942%</b>	3850%
Double Dunk	545%	664%	595%	<b>1686%</b>	1064%
Enduro	85%	157%	141%	243%	<b>250%</b>
Fishing Derby	164%	213%	202%	<b>248%</b>	230%
Freeway	104%	111%	112%	<b>114%</b>	111%
Frostbite	17%	9%	38%	<b>101%</b>	79%
Gopher	395%	820%	677%	1496%	<b>2266%</b>
Gravitar	9%	11%	8%	<b>12%</b>	5%
H.E.R.O.	65%	66%	64%	74%	<b>88%</b>
Ice Hockey	77%	76%	70%	<b>103%</b>	95%
James Bond 007	270%	539%	485%	<b>1870%</b>	1188%
Kangaroo	242%	431%	434%	<b>541%</b>	484%
Krull	639%	<b>1051%</b>	592%	762%	812%
Kung-Fu Master	115%	115%	131%	175%	<b>192%</b>
Montezuma's Revenge	<b>0%</b>	<b>0%</b>	<b>0%</b>	<b>0%</b>	<b>0%</b>
Ms. Pac-Man	42%	47%	36%	<b>93%</b>	67%
Name This Game	103%	143%	145%	173%	<b>194%</b>
Phoenix	119%	179%	177%	281%	<b>494%</b>
Pitfall!	-1%	2%	3%	-2%	<b>3%</b>
Pong	114%	<b>118%</b>	118%	117%	117%
Private Eye	0%	<b>5%</b>	0%	0%	0%
Q*bert	97%	115%	112%	121%	<b>140%</b>
River Raid	38%	83%	86%	84%	<b>121%</b>
Road Runner	505%	794%	563%	735%	<b>801%</b>
Robotank	636%	561%	<b>648%</b>	623%	581%
Seaquest	14%	31%	39%	63%	<b>106%</b>
Skiing	32%	-101%	<b>63%</b>	56%	56%
Solaris	20%	-9%	17%	<b>28%</b>	4%
Space Invaders	102%	144%	156%	179%	<b>496%</b>
Stargunner	559%	619%	620%	<b>653%</b>	584%
Surround	27%	22%	43%	<b>115%</b>	73%
Tennis	<b>232%</b>	147%	6%	154%	154%
Time Pilot	78%	231%	287%	339%	<b>474%</b>
Tutankham	36%	102%	<b>133%</b>	124%	49%
Up'n Down	85%	160%	201%	140%	<b>301%</b>
Venture	14%	6%	8%	5%	<b>73%</b>
Video Pinball	12793%	16342%	20814%	18834%	<b>27652%</b>
Wizard of Wor	51%	47%	165%	101%	<b>234%</b>
Yars' Revenge	29%	<b>35%</b>	17%	16%	26%
Zaxxon	58%	73%	111%	114%	<b>147%</b>

Table 7: Same as Table 6 but using no-ops evaluation.

Game			DQN			Double DQN		
	random	human	baseline	Gorila	rank-b.	baseline	rank-b.	prop.
Alien	128.3	6371.3	634.0	813.5	1191.0	1033.4	<b>1334.7</b>	900.5
Amidar	11.8	1540.4	178.4	189.2	98.9	165.8	129.1	<b>218.4</b>
Assault	166.9	628.9	1672.5	1195.8	3081.3	5934.5	6548.9	<b>7748.5</b>
Asterix	164.5	7536.0	3170.5	3324.7	9199.5	16837.0	22484.5	<b>31907.5</b>
Asteroids	871.3	36517.3	1458.7	933.6	1677.2	1189.3	<b>1745.1</b>	1654.0
Atlantis	13463.0	26575.0	282096.0	<b>629166.5</b>	207526.0	207658.0	330647.0	593642.0
Bank Heist	21.7	644.5	261.5	399.4	823.7	751.3	<b>876.6</b>	816.8
Battlezone	3560.0	33030.0	23750.0	19938.0	22250.0	24740.0	25520.0	<b>29100.0</b>
Beam Rider	254.6	14961.0	9743.2	3822.1	12041.9	17417.2	<b>31181.3</b>	26172.7
Berzerk	196.1	2237.5	493.4		644.0	990.7	865.9	<b>1165.6</b>
Bowling	35.2	146.5	56.5	54.0	58.0	<b>69.6</b>	52.0	65.8
Boxing	-1.5	9.6	70.3	<b>74.2</b>	69.6	73.5	72.3	68.6
Breakout	1.6	27.9	354.5	313.0	<b>481.1</b>	368.9	343.0	371.6
Centipede	1925.5	10321.9	3973.9	<b>6296.9</b>	2959.4	3765.7	3489.1	3421.9
Chopper Command	644.0	8930.0	5017.0	3191.8	<b>6685.0</b>	2919.0	4635.0	6604.0
Crazy Climber	9337.0	32667.0	98128.0	65451.0	109337.0	100884.0	127512.0	<b>131086.0</b>
Defender	1965.5	14296.0	15917.5		20634.0	<b>27510.0</b>	23666.5	21093.5
Demon Attack	208.3	3442.8	12550.7	14880.1	19478.8	55136.8	61277.5	<b>73185.8</b>
Double Dunk	-16.0	-14.4	-6.0	-11.3	-5.3	-3.5	<b>16.0</b>	2.7
Enduro	-81.8	740.2	626.7	71.0	1265.6	1182.4	1831.0	<b>1884.4</b>
Fishing Derby	-77.1	5.1	-16.1	4.6	3.5	1.8	<b>9.8</b>	9.2
Freeway	0.1	25.6	26.9	10.2	28.4	28.8	<b>28.9</b>	27.9
Frostbite	66.4	4202.8	496.1	426.6	288.7	1448.1	<b>3510.0</b>	2930.2
Gopher	250.0	2311.0	7936.0	4373.0	17478.2	15253.0	34858.8	<b>57783.8</b>
Gravitar	245.5	3116.0	298.0	<b>538.4</b>	351.0	200.5	269.5	218.0
H.E.R.O.	1580.3	25839.4	13748.7	8963.4	15150.9	14892.5	<b>20889.9</b>	20506.4
Ice Hockey	-9.7	0.5	-1.6	-1.7	-3.8	-2.5	<b>-0.2</b>	-1.0
James Bond 007	33.5	368.5	697.5	444.0	1074.5	573.0	<b>3961.0</b>	3511.5
Kangaroo	100.0	2739.0	4496.0	1431.0	9053.0	11089.0	<b>12185.0</b>	10241.0
Krull	1151.9	2109.1	6206.0	6363.1	<b>11209.5</b>	6796.1	6872.8	7406.5
Kung-Fu Master	304.0	20786.8	20882.0	20620.0	20181.0	30207.0	<b>31676.0</b>	31244.0
Montezuma's Revenge	25.0	4182.0	47.0	<b>84.0</b>	44.0	42.0	51.0	13.0
Ms. Pac-Man	197.8	15375.0	980.2	1263.0	964.7	907.5	<b>1865.9</b>	1824.6
Name This Game	1747.8	6796.0	6738.8	9238.5	8738.5	8919.7	10497.6	<b>11836.1</b>
Phoenix	1134.4	6686.2	7484.8		16107.8	12366.5	16903.6	<b>27430.1</b>
Pitfall!	-348.8	5998.9	-878.3		-193.7	-186.7	-427.0	<b>-14.8</b>
Pong	-18.0	15.5	18.0	16.7	18.7	18.8	<b>18.9</b>	<b>18.9</b>
Private Eye	662.8	64169.1	207.9	<b>2598.6</b>	2202.3	-683.8	670.7	179.0
Q*bert	183.0	12085.0	9271.5	7089.8	<b>12740.5</b>	10545.5	9944.0	11277.0
River Raid	588.3	14382.2	4748.5	5310.3	10205.5	10759.9	11807.2	<b>18184.4</b>
Road Runner	200.0	6878.0	35215.0	43079.8	<b>57207.0</b>	43156.0	52264.0	56990.0
Robotank	2.4	8.9	58.7	<b>61.8</b>	51.3	59.1	56.2	55.4
Seaquest	215.5	40425.8	4120.5	10145.9	11848.8	14498.0	25463.7	<b>39096.7</b>
Skiing	-15287.4	-3686.6	-12142.1		-29404.3	-11490.4	<b>-10169.1</b>	-10852.8
Solaris	2047.2	11032.6	1295.4		134.6	810.0	<b>2272.8</b>	2238.2
Space Invaders	182.6	1464.9	1293.8	1183.3	1696.9	2628.7	3912.1	<b>9063.0</b>
Stargunner	697.0	9528.0	52970.0	14919.2	58946.0	58365.0	<b>61582.0</b>	51959.0
Surround	-9.7	5.4	-6.0		-5.3	1.6	<b>5.9</b>	-0.9
Tennis	-21.4	-6.7	<b>11.1</b>	-0.7	-2.3	-7.8	-5.3	-2.0
Time Pilot	3273.0	5650.0	4786.0	<b>8267.8</b>	5391.0	6608.0	5963.0	7448.0
Tutankham	12.7	138.3	45.6	<b>118.5</b>	96.5	92.2	56.9	33.6
Up'n Down	707.2	9896.1	8038.5	8747.7	16626.5	19086.9	12157.4	<b>29443.7</b>
Venture	18.0	1039.0	136.0	<b>523.4</b>	110.0	21.0	94.0	244.0
Video Pinball	20452.0	15641.1	154414.1	112093.4	214925.3	367823.7	295972.8	<b>374886.9</b>
Wizard of Wor	804.0	4556.0	1609.0	<b>10431.0</b>	2755.0	6201.0	5727.0	7451.0
Yars' Revenge	1476.9	47135.2	2694.3		<b>6626.7</b>	6260.6	4687.4	5965.1
Zaxxon	475.0	8443.0	4412.0	6159.4	5901.0	8504.0	9474.0	<b>9501.0</b>

Table 8: Scores obtained on the original 49 Atari games plus 8 others with human starts evaluation. Also includes the random and human baseline scores used for normalization. Human and random scores are taken from Nair et al. (2015), with scores for the additional 8 games. DQN and Double DQN scores are from van Hasselt et al. (2016). Note that the Gorila results used much more data and computation, but the other methods are more directly comparable to each other in this respect.

Game	DQN				Double DQN		
	random	human	baseline	rank-b.	baseline	rank-b	prop.
Alien	227.8	7127.7	1620.0	4860.1	3747.7	4203.8	<b>6648.6</b>
Amidar	5.8	1719.5	978.0	812.8	1793.3	1838.9	<b>2051.8</b>
Assault	222.4	742.0	4280.4	3438.2	5393.2	7672.1	<b>7965.7</b>
Asterix	210.0	8503.3	4359.0	10681.0	17356.5	31527.0	<b>41268.0</b>
Asteroids	719.1	47388.7	1364.5	1762.9	734.7	<b>2654.3</b>	1699.3
Atlantis	12850.0	29028.1	279987.0	231403.0	106056.0	357324.0	<b>427658.0</b>
Bank Heist	14.2	753.1	455.0	1121.4	1030.6	1054.6	<b>1126.8</b>
Battlezone	2360.0	37187.5	29900.0	23880.0	31700.0	31530.0	<b>38130.0</b>
Beam Rider	363.9	16926.5	8627.5	10312.8	13772.8	<b>23384.2</b>	22430.7
Berzerk	123.7	2630.4	585.6	844.5	1225.4	1305.6	<b>1614.2</b>
Bowling	23.1	160.7	50.4	48.0	<b>68.1</b>	47.9	62.6
Boxing	0.1	12.1	88.0	<b>99.5</b>	91.6	95.6	98.8
Breakout	1.7	30.5	385.5	<b>422.9</b>	418.5	373.9	381.5
Centipede	2090.9	12017.0	4657.7	3401.5	<b>5409.4</b>	4463.2	5175.4
Chopper Command	811.0	7387.8	6126.0	6649.0	5809.0	<b>8600.0</b>	5135.0
Crazy Climber	10780.5	35829.4	110763.0	123221.0	117282.0	141161.0	<b>183137.0</b>
Defender	2874.5	18688.9	23633.0	25248.5	<b>35338.5</b>	31286.5	24162.5
Demon Attack	152.1	1971.0	12149.4	21318.2	58044.2	<b>71846.4</b>	70171.8
Double Dunk	-18.6	-16.4	-6.6	-4.0	-5.5	<b>18.5</b>	4.8
Enduro	0.0	860.5	729.0	1353.2	1211.8	2093.0	<b>2155.0</b>
Fishing Derby	-91.7	-38.7	-4.9	21.2	15.5	<b>39.5</b>	30.2
Freeway	0.0	29.6	30.8	33.0	33.3	<b>33.7</b>	32.9
Frostbite	65.2	4334.7	797.4	441.2	1683.3	<b>4380.1</b>	3421.6
Gopher	257.6	2412.5	8777.4	17921.2	14840.8	32487.2	<b>49097.4</b>
Gravitar	173.0	3351.4	473.0	523.0	412.0	<b>548.5</b>	330.5
H.E.R.O.	1027.0	30826.4	20437.8	20547.5	20130.2	23037.7	<b>27153.9</b>
Ice Hockey	-11.2	0.9	-1.9	-2.0	-2.7	<b>1.3</b>	0.3
James Bond 007	29.0	302.8	768.5	1504.5	1358.0	<b>5148.0</b>	3283.0
Kangaroo	52.0	3035.0	7259.0	12902.0	12992.0	<b>16200.0</b>	14492.0
Krull	1598.0	2665.5	8422.3	<b>12815.9</b>	7920.5	9728.0	10263.1
Kung-Fu Master	258.5	22736.3	26059.0	26047.0	29710.0	39581.0	<b>43470.0</b>
Montezuma's Revenge	0.0	4753.3	<b>0.0</b>	<b>0.0</b>	<b>0.0</b>	<b>0.0</b>	<b>0.0</b>
Ms. Pac-Man	307.3	6951.6	3085.6	3414.3	2711.4	<b>6518.7</b>	4751.2
Name This Game	2292.3	8049.0	8207.8	10497.2	10616.0	12270.5	<b>13439.4</b>
Phoenix	761.4	7242.6	8485.2	12353.6	12252.5	18992.7	<b>32808.3</b>
Pitfall!	-229.4	6463.7	-286.1	-107.4	-29.9	-356.5	<b>0.0</b>
Pong	-20.7	14.6	19.5	<b>21.0</b>	20.9	20.6	20.7
Private Eye	24.9	69571.3	146.7	<b>3441.9</b>	129.7	200.0	200.0
Q*bert	163.9	13455.0	13117.3	15470.0	15088.5	16256.5	<b>18802.8</b>
River Raid	1338.5	17118.0	7377.6	14494.6	14884.5	14522.3	<b>20494.0</b>
Road Runner	11.5	7845.0	39544.0	62234.0	44127.0	57608.0	<b>62785.0</b>
Robotank	2.2	11.9	63.9	56.6	<b>65.1</b>	62.6	58.6
Seaquest	68.4	42054.7	5860.6	13051.9	16452.7	26357.8	<b>44417.4</b>
Skiing	-17098.1	-4336.9	-13062.3	-29975.8	<b>-9021.8</b>	-9996.9	-9900.5
Solaris	1236.3	12326.7	3482.8	208.0	3067.8	<b>4309.0</b>	1710.8
Space Invaders	148.0	1668.7	1692.3	2340.8	2525.5	2865.8	<b>7696.9</b>
Stargunner	664.0	10250.0	54282.0	59978.0	60142.0	<b>63302.0</b>	56641.0
Surround	-10.0	6.5	-5.6	-6.3	-2.9	<b>8.9</b>	2.1
Tennis	-23.8	-8.3	<b>12.2</b>	-1.0	-22.8	0.0	0.0
Time Pilot	3568.0	5229.2	4870.0	7398.0	8339.0	9197.0	<b>11448.0</b>
Tutankham	11.4	167.6	68.1	171.0	<b>218.4</b>	204.6	87.2
Up'n Down	533.4	11693.2	9989.9	18370.6	22972.2	16154.1	<b>34082.7</b>
Venture	0.0	1187.5	163.0	69.0	98.0	54.0	<b>863.0</b>
Video Pinball	16256.9	17667.9	196760.4	246844.3	309941.9	282007.3	<b>406420.4</b>
Wizard of Wor	563.5	4756.5	2704.0	2533.0	7492.0	4802.0	<b>10373.0</b>
Yars' Revenge	3092.9	54576.9	18098.9	<b>21151.0</b>	11712.6	11357.0	16451.7
Zaxxon	32.5	9173.3	5363.0	6715.0	10163.0	10469.0	<b>13490.0</b>

Table 9: Scores obtained on 57 Atari games, as in Table 8 but with no-ops evaluation.