
Low-Precision Reinforcement Learning: Running Soft Actor-Critic in Half Precision

Johan Bjorck¹ Xiangyu Chen¹ Christopher De Sa¹ Carla P. Gomes¹ Kilian Q. Weinberger¹

Abstract

Low-precision training has become a popular approach to reduce compute requirements, memory footprint, and energy consumption in supervised learning. In contrast, this promising approach has not yet enjoyed similarly widespread adoption within the reinforcement learning (RL) community, partly because RL agents can be notoriously hard to train even in full precision. In this paper we consider continuous control with the state-of-the-art SAC agent and demonstrate that a naïve adaptation of low-precision methods from supervised learning fails. We propose a set of six modifications, all straightforward to implement, that leaves the underlying agent and its hyperparameters unchanged but improves the numerical stability dramatically. The resulting modified SAC agent has lower memory and compute requirements while matching full-precision rewards, demonstrating that low-precision training can substantially accelerate state-of-the-art RL without parameter tuning.

1. Introduction

Reinforcement learning (RL) is a promising paradigm for constructing autonomous agents for transportation (Kendall et al., 2019), robotics (Zhu et al., 2020), scientific applications (Bellemare et al., 2020) and beyond (Silver et al., 2016). To enable such impressive feats, RL often requires large amounts of computation and samples from the underlying environments (Silver et al., 2016; Jaderberg et al., 2019; Berner et al., 2019). Consequently, many approaches for scaling RL have been proposed, including distributed (Horgan et al., 2018; Wijmans et al., 2019) and asynchronous training (Mnih et al., 2016), GPU-accelerated environments (Dalton et al., 2019) and many others (Petrov et al., 2020; Liu et al., 2020a). Within supervised

learning, an emerging trend for accelerating deep learning is low-precision training, where one uses fewer than 32 bits to represent individual parameters, activations, and gradients (Gupta et al., 2015). This speeds up computations and decreases latency, memory footprint, and energy consumption (De Sa et al., 2017; Reuther et al., 2019; Gong et al., 2018). Such benefits can accelerate RL research, allow longer deployment before recharging batteries (Tomy et al., 2019), lower latency in high-speed applications (Liu et al., 2020b), and is especially relevant in RL from raw pixel observations which can be compute-intensive (Zhan et al., 2020). Despite such upsides, the RL community has not adopted low-precision training, in part as RL agents are notoriously hard to train even in full precision (Henderson et al., 2017; Engstrom et al., 2020; Jordan et al., 2020).

The goal of this paper is to demonstrate the feasibility of RL in low precision without changing hyperparameters or the underlying RL algorithm. Specifically, we focus on continuous control environments (Tassa et al., 2020) and consider the state-of-the-art soft actor-critic (SAC) algorithm (Haarnoja et al., 2018). We demonstrate that strategies developed for supervised learning, such as loss scaling and mixed-precision (Micikevicius et al., 2017), fail in this context. To enable RL in low precision, we propose a set of six methods that improve the numerical stability without changing the underlying agent. Some of these are novel, for example, we propose to store the square root of the second moment in Adam to decrease the needed dynamic range. Others have more of an engineering flavor, e.g. reordering arithmetic operations, but are nonetheless crucially needed for performant agents.

We experimentally verify that with these methods it is possible to train SAC agents in low precision with comparable performance to full-precision agents, thus demonstrating the feasibility of low-precision RL. We further benchmark compute time and memory consumption and find dramatic improvements in both aspects. We perform experiments on RL from both environment state representations and pixel observations, and also perform ablation experiments and simulate various numerical formats with torch (Zhang et al., 2019). Finally, we release our code to encourage future work on RL in low precision. Debugging numerical

¹Department of Computer Science, Cornell University, USA. Correspondence to: Johan Bjorck <njb225@cornell.edu>.

issues can be problematic as they do not necessarily result in crashes. We hope that the techniques and code we introduce can provide a starting point for other researchers interested in low-precision RL.

2. Background

2.1. Low-Precision Deep Learning

Within deep learning, model parameters are typically represented as 32-bit floating-point numbers (fp32) (Kahan, 1996). One bit is used for the sign, the exponent uses 8 bits and the rest represent the significand. If we interpret these bits as integers s, e, m , the value they represent is

$$\text{value}(s, e, m) = (-1)^s \times 2^{e-127} \times 1.m$$

Arithmetic in low precision is precarious. With a finite number of bits, only a finite set of numbers can be represented. Calculations yielding numbers too large to be represented cause *overflow* and result in ∞ (which has a specific bit-pattern). Conversely, numbers too small to be represented cause *underflow* and are cast to 0. If a is small and b is large, in low precision the representable number closest to $a + b$ might be b . This can effectively undo addition. We refer to such issues as *precision* issues. Finally, there is also a bit pattern for *NaN* (abbreviation for "not a number") which can come from division by 0 or arithmetic with ∞ .

Low-precision deep learning relies on using fewer than 32 bits for individual parameters, activations, and gradients (Gupta et al., 2015), yielding faster computation, lower latency, less power consumption, and smaller memory footprint (De Sa et al., 2017; Gong et al., 2018). Nvidia GPUs (Whitehead & Fit-Florea, 2011) and PyTorch (Paszke et al., 2019) currently support 16-bit floating-point numbers (fp16 or half-precision), which uses 5 bits in the exponent (Kahan, 1996). Other representations are actively being researched (Kalamkar et al., 2019; Das et al., 2018; Song et al., 2017). For low-precision training of neural networks, a few common tricks are used to improve numerical stability. One can scale the loss to prevent gradient underflow in the backward pass, or mix high and low precision (Mickevičius et al., 2017), or coerce Nan values to 0 and ∞ to some large number. In Figure 1 we apply such baseline methods to our RL setting – the SAC agent evaluated on the planet benchmark (Hafner et al., 2019) – and observe that they do not reach competitive performance (see Section 4 for details.) This suggests that novel strategies are needed to enable low-precision RL.

2.2. Reinforcement Learning

Continuous control tasks, e.g., robotic arm control, are often formulated as Markov decision processes (MDPs) defined by a tuple $(\mathcal{S}, \mathcal{A}, P, r)$ (Sutton & Barto, 2018). The state




Figure 1. Returns for the SAC agent, averaged across tasks in the planet-benchmarks, when using low-precision methods from supervised learning (Mickevičius et al., 2017). Naively adopting these methods in RL gives poor performance, suggesting that novel methods are needed. Fp16 always crashes, yielding no rewards.

space \mathcal{S} and action space \mathcal{A} are both continuous and might be bounded. At each timestep t , the agent is in a state $s_t \in \mathcal{S}$ and takes an action $a \in \mathcal{A}$ to arrive at a new state $s_{t+1} \in \mathcal{S}$. The transition between states given an action is random with transition probability $P : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \rightarrow [0, \infty)$. Furthermore, at each timestep, the agent receives a reward r_t as per the reward distribution r . The agent is typically trained to take actions that maximize expected discounted reward $\mathbb{E}[\sum_i \gamma^i r_i]$ for some fixed $0 < \gamma < 1$ which defines a horizon for the problem. The performance is typically measured by the cumulative rewards $\sum_i r_i$.

2.3. Soft Actor-Critic (SAC)

Soft actor-critic (SAC) is a popular RL algorithm for continuous control (Haarnoja et al., 2018), forming the basis for recent state-of-the-art algorithms such as CURL (Laskin et al., 2020b), RAD (Laskin et al., 2020a) and DRQ (Kostrikov et al., 2020). For every state s , SAC defines a continuous multivariate *policy* distribution $\pi_\theta(s)$ over actions, defined via a neural network with parameters θ . The neural network outputs two vectors, a mean $\mu_\theta(s)$ and a standard deviation $\sigma_\theta(s)$ and the policy is defined as

$$\pi_\theta(s) = \tanh(\mu_\theta(s) + \epsilon \odot \sigma_\theta(s)), \quad \epsilon \sim \mathcal{N}(0, 1). \quad (1)$$

Here, \tanh is applied elementwise, yielding the action space $[-1, 1]^n$ for some $n = \text{dim}(\mathcal{A})$. An action a can then be sampled from $\pi_\theta(s)$ simply by sampling ϵ from the standard normal distribution. To encourage *exploration*, the entropy $H(\pi_\theta(s_t))$ of the action distribution (times a hyperparameter α) is added to the reward at each state s_t . In practice, the entropy can be calculated by sampling an action a from π_θ and thereafter calculating $\log \pi_\theta(a|s)$.

The agent also uses a neural network with parameters ψ , of-

ten referred to as the *critic*, that outputs a q-value $Q_\psi(\mathbf{s}, \mathbf{a})$ (Watkins, 1989) for each state \mathbf{s} and action \mathbf{a} . Q-values roughly measure the expected rewards that can be obtained by taking action \mathbf{a} in state \mathbf{s} . The critic network is trained to minimize the square loss between its output $Q_\psi(\mathbf{a}, \mathbf{s}_t)$ and the target $r_t + \gamma \mathbb{E}[\hat{Q}(\mathbf{a}_{t+1}, \mathbf{s}_{t+1}) + \alpha H(\pi)]$. The value $\hat{Q}(\mathbf{s}, \mathbf{a})$ is obtained from another network – the *target* network – which shares architecture with the critic but whose weights $\hat{\psi}$ are the exponentially averaged weights of the critic network, i.e., we do not take gradients w.r.t $\hat{\psi}$ and for some β we occasionally update

$$\hat{\psi} \leftarrow \beta \hat{\psi} + (1 - \beta)\psi.$$

The policy π_θ is optimized to maximize the expected rewards. This can be done by sampling an action \mathbf{a} from $\pi_\theta(\mathbf{s}_t)$ and thereafter taking derivatives w.r.t. \mathbf{a} through the q-values $Q_\psi(\mathbf{s}, \mathbf{a})$. In practice, instead of treating α as a fixed hyperparameter, one can update α with a gradient descent type optimizer so that the average entropy of states matches some predefined value. SAC is typically optimized by the Adam optimizer (Kingma & Ba, 2014). For further details, see Haarnoja et al. (2018).

Table 1. Proposed modifications and the problems they solve.

#	method	problem
1	hAdam	under/overflow
2	softplus-fix	overflow
3	normal-fix	underflow
4	Kahan-momentum	precision/underflow
5	compound loss scaling	underflow
6	Kahan-gradients	precision

3. Proposed Methods

To improve the numerical stability of SAC (Haarnoja et al., 2018) we propose a set of six modifications that do not change the underlying algorithm. These are listed in Table 1. We point out that some of them are relatively straightforward engineering choices; nonetheless, such implementation level modifications are often important to achieve good performance in RL (Engstrom et al., 2020). For the sake of being self-contained, we present all modifications here. We verify that these components individually contribute to improved performance in Figure 3. However, not all these components are needed for all environments – we have found that some environments are less sensitive to numerical issues; see appendix E for details. Many of these methods are agnostic to the underlying RL method and we hope that they provide a sensible starting point for other RL agents.

1. Storing the Hypotenuse in Adam. The Adam optimizer stores the exponential moving average of the squared gradient. It is stored in a variable v which is updated as $v_{t+1} \leftarrow \beta_2 v_t + (1 - \beta_2) g_{t+1}^2$ for gradient g_{t+1} . Storing the

full range of v is infeasible in low precision; for example, if $g_{t+1} = 10^{-7}$ computing g_{t+1}^2 might cause underflow in 16-bit precision. We instead propose to store $w = \sqrt{v}$ which has a smaller dynamic range than v . To update w we use the hypot function, which is defined as

$$\text{hypot}(a, b) = \sqrt{a^2 + b^2}.$$

In a naïve implementation, the intermediate results a^2 or b^2 might underflow. To avoid this, we can, e.g., rewrite the expression above as follows for nonzero a and b :

$$\text{hypot}(a, b) = \max(a, b) \sqrt{1 + \left(\frac{\min(a, b)}{\max(a, b)}\right)^2}.$$

This expression is more numerically stable when a and b can be represented, but a^2 and b^2 are rounded to 0. To allow for $a = 0, b = 0$ one can add a numerical ϵ to the denominator. Once we have a numerically suitable hypot function we define the updates for w as follows

$$w_{t+1} \leftarrow \text{hypot}(\sqrt{\beta_2} w_t, \sqrt{(1 - \beta_2)} g_{t+1}).$$

Under this update rule w retains the semantics $w = \sqrt{v}$. We update the network parameters θ with $\theta \leftarrow \theta - \alpha \frac{m_t}{w_t + \epsilon}$ for some learning rate α and the ϵ used in Adam. In practice, this exchanges the square root computation to the more numerically stable hypot execution. Both have comparable complexity as they require finding a square root via iterative methods. In fact, hypot is a common numerical primitive (found in, e.g., <math.h> in the C language). We can pre-compute the values of $\sqrt{\beta_2}$ and $\sqrt{(1 - \beta_2)}$ up-front. Although this modification might cause a slight computational overhead, depending on the hypot implementation, the gradient computation dominates the computational cost. The resulting method, which we call **hAdam**, is described in Algorithm 1.

Algorithm 1 hAdam.

```

m ← 0
w ← 0
for t = 0, 1, 2, ... do
    g ← ∇θ ℓ
    m ← β1 m + (1 - β1) g
    w ← hypot(√β2 w, √(1 - β2) g)
    m ← m / (1 - β1t) // bias correction
    w ← w / √(1 - β2t) // bias correction
    θ ← θ - α m / (w + ε)

```

end

2 & 3. Numerical Issues in the Policy. Naïvely handling the policy parametrization (eq. (1)) of SAC leads to poor performance in low precision. Calculating the log-probability of actions, which is needed to compute the entropy, is problematic in low precision. Recall that actions \mathbf{a} are drawn

from $\tanh(\mathbf{u})$ where $\mathbf{u} = \boldsymbol{\mu}_\theta + \boldsymbol{\epsilon} \odot \boldsymbol{\sigma}_\theta$ and $\boldsymbol{\epsilon} \sim \mathcal{N}(0, 1)$. Here, $\boldsymbol{\mu}_\theta$ and $\boldsymbol{\sigma}_\theta$ are the outputs from the policy neural network. To calculate the log-probability, we need to account for the change of variables due to the tanh transformation:

$$\begin{aligned} \log \pi(\mathbf{a}|\mathbf{s}) &= \log P(\mathbf{u}|\mathbf{s}) - \sum_i \log(1 - \tanh^2(u_i)) \\ &= \log P(\mathbf{u}|\mathbf{s}) - \sum_i 2[\log(2) - u - \log(1 + \exp(-2u_i))]. \end{aligned}$$

The first expression is numerically unsuitable as $\tanh^2(u_i)$ can be rounded to 1 with insufficient precision. The latter expression is known to be more numerically stable (Kostrikov et al., 2020), but can still overflow in the backward pass when $\exp(-2u_i)$ is large. This happens, e.g., in PyTorch (Paszke et al., 2019). To prevent such overflow, we exchange $\log(1 + \exp(-2u_i))$ by a linear function (which has a stable backward pass) for sufficiently large negative u_i , i.e when $u < K$ for a K that is chosen depending on the dynamic range of the number representation. The resulting function is shown in eq. (2). This modification only carries a minimal computational overhead and we call it **softplus-fix**.

$$\text{softplus}'(u_i) = \begin{cases} -2u_i & \text{if } u < K \\ \log(1 + \exp(-2u_i)) & \text{else} \end{cases} \quad (2)$$

Secondly, we need to calculate $\log P(\mathbf{u}|\mathbf{s})$ which is just the log-probability of a normal variable. The numerical stability of this operation is important, and, e.g., the PyTorch implementation of the normal distribution has log-probabilities calculated as $\log x \sim \frac{(x-\mu)^2}{\sigma^2}$ (we omit normalization constants here). If x, μ, σ are all very small, it is possible that, e.g., σ^2 is rounded to zero in 16-bit precision despite the correct ratio being around 1. To improve the stability, we instead calculate the log-probability as $\log x \sim \left(\frac{x-\mu}{\sigma}\right)^2$ (modulo normalization constant). This turns out to solve such underflow issues for the SAC implementation of Yarats & Kostrikov (2020) and incurs no computational or memory cost. We call this modification the **normal-fix**.

4. Stability of Soft-Updates. Recall that the target network is an exponential moving average of the critic network:

$$\hat{\psi} \leftarrow \beta \hat{\psi} + (1 - \beta)\psi.$$

For β close to one, $(1 - \beta)\psi$ can be very small which can cause numerical issues. To compute the momentum update of the target network, we propose to use Kahan summation (Kahan, 1965), a general method for calculating sums of low-precision numbers. When summing a long sequence of numbers in low precision, the numerical error can grow with the length of the sequence. Kahan summation adds a compensation variable that is updated as each term of the sum is added, which improves numerical stability. The

method is given in algorithm 2. We rewrite the momentum update as adding $(1 - \beta)(\psi - \hat{\psi})$ to $\hat{\psi}$ and then use Kahan summation to add this term to $\hat{\psi}$. To prevent underflow of $(1 - \beta)(\psi - \hat{\psi})$ we can scale it by some constant $C > 1$ and add them to a copy of $\hat{\psi}$ where all weights are scaled by C . This strategy has a relatively small computational cost but larger memory costs, however, this is offset by storing the model in lower precision. Additionally, we still reap the benefits of computing gradients and activations in lower-precision, which often is the main source of memory consumption. We call this method **Kahan-momentum**.

Algorithm 2 Kahan summation (Kahan, 1965).

```

sum ← 0
c ← 0
for val in values do
    y ← val - c
    t ← sum + y
    c ← (t - sum) - y
    sum ← t
end

```

end

5. Compound Loss Scaling. Adam stores the exponential moving averages of the gradients and the squared gradients, updating them as, e.g., $m \leftarrow \beta_1 m + (1 - \beta_1)g$. For small gradients g and β_1 close to 1, the term $(1 - \beta_1)g$ might be rounded to zero. Fortunately, the Adam update is invariant if we rescale the numerator and the denominator simultaneously. Thus, we can choose some fixed large γ and update the Adam buffers as follows

$$\begin{aligned} m &\leftarrow \beta_1 m + (1 - \beta_1)\gamma g, \\ w &\leftarrow \text{hypot}(\sqrt{\beta}w, \sqrt{(1 - \beta)\gamma\gamma}). \end{aligned}$$

This modifications ensure that $(1 - \beta_1)\gamma g$ does not underflow. To retain the semantics of Adam, we update the parameters as $\theta \leftarrow \theta - \alpha \frac{m}{w + \gamma\epsilon}$. For implementing this strategy we scale the loss by some γ – this scales the gradients and thus the contents of m and w . Due to its similarity with loss scaling, we call this method **compound loss scaling**. Compared to loss scaling we do not need to "unscale" the gradients. As is standard in loss scaling, we change γ dynamically following the strategy of PyTorch amp (Carilli, 2020); see appendix B for details. This results in a small extra computational cost as we need to check that the gradients do not overflow when dynamically updating γ .

6. Kahan Summation of Gradients. Lastly, we perform Kahan summation on the gradient updates for the weights ψ of the critic network and for the parameter α . This trick turns out not to be needed for the actor-network. Formally, for a gradient update $\theta \leftarrow \theta + \Delta\theta$, we apply algorithm 2. Again, this method comes with a small computational overhead but needs a memory footprint comparable to that of the actual

model. We already save some memory from storing the parameters in low precision, and the main source of memory consumption is often computing the gradients, which is not affected. We call this method **Kahan-gradients**.

3.1. Guarantees

In practice, our modifications improve numerical stability. However, one might be worried that the underlying RL algorithm is changed. As the methods rely on algebraically rewriting operations, the underlying algorithm will be the same – at least in infinite numerical precision where arithmetic is commutative and associative. We formalize this intuition as follows; see Appendix C for a proof.

Statement 1 *In infinite precision, training SAC with the modification (normal-fix, softplus-fix, hAdam, Kahan-momentum, compound loss scaling, and Kahan-gradients) is equivalent to training without them.*

4. Experiments

4.1. Setup

To verify the efficacy of our proposed methods we now perform experiments on the state-of-the-art agent SAC (Haarnoja et al., 2018). We here focus on three specific questions about our method i) does it match rewards

from fp32 training? ii) does it improve compute/memory consumption? iii) are all proposed components needed? We start from the publicly accessible SAC implementation of Yarats & Kostrikov (2020) which is used in recent state-of-the-art algorithms such as CURL (Laskin et al., 2020b), RAD (Laskin et al., 2020a) and DRQ (Kostrikov et al., 2020). For environments, we consider the planet benchmark, popularized by (Hafner et al., 2019) and used in e.g. Kostrikov et al. (2020); Laskin et al. (2020b;a). It consists of six continuous control tasks from the deep mind control suite (Tassa et al., 2020): finger spin, cartpole swingup, reacher easy, cheetah run, walker walk, and ball in cup catch. We do not tune hyperparameters, but instead use hyperparameters from Yarats & Kostrikov (2020) (listed in Appendix B). We measure performance after 500,000 environment steps as per (Kostrikov et al., 2020; Laskin et al., 2020b; Lee et al., 2019) with individual runs scored by the average return over 10 episodes. We use 15 seeds for the main experiments (Figure 2 and Figure 5) and 5 seeds for all other experiments. Runs using SAC naïvely ported to fp16 can crash when taking actions that are non-finite, these are scored as 0. Figures show means and standard deviations (stds) across runs; when aggregating stds across tasks, we calculate the stds of individual tasks and then plot the average std. We primarily perform experiments with fp16 as this format is natively supported in PyTorch and on Nvidia V100 GPUs. Fp16 has both a smaller dynamic range




Figure 2. Learning curves for SAC (Haarnoja et al., 2018) on the planet-benchmark for standard fp32 training and training in fp16 with our proposed methods. Naive fp16 methods fail as per Figure 1. Performance is very close for all environments, demonstrating that low-precision RL is possible without changing the underlying RL algorithm or its hyperparameters.

and lower precision compared to fp32. We simulate other numerical formats in section 4.5.

4.2. Comparison to Fp32

We first compare the performance of the SAC agent trained in fp32 and fp16, the latter using all our methods from section 3. Learning curves for different tasks are given in Figure 2. We see that half-precision training performs very close to full-precision training, validating that our methods essentially match fp32 performance despite training in half-precision.

4.3. Comparison to Baselines

We now compare our strategy to the following standard low-precision methods from supervised learning:

- Numeric coercion, abbreviated as *coerc*. NaNs are coerced into zero and $\pm\infty$ to the largest/smallest fp16 value.
- Loss scaling (Micikevicius et al., 2017), abbreviated as *loss scale*. We perform loss scaling with dynamic scale, using the same parameters as for our compound scaling. In appendix E we use the *amp* (Carilli, 2020) default parameters which gives similar results.
- Mixed precision and loss scaling (Micikevicius et al., 2017), abbreviated as *mixed precision*. In addition to loss scaling, we use mixed precision. Gradients, Adam buffers, and parameters are stored in fp32. Parameters are cast to half-precision in the forward pass, and the backward pass is calculated in half-precision.

In Figure 1 we show the performance of these baselines from supervised learning. These methods perform poorly in low-precision RL, demonstrating that additional strategies are necessary in this context.

4.4. Ablation Experiments

We additionally perform an ablation experiment to ensure that all our proposed methods improve performance individually. The experiment is structured as follows: we start with the original SAC agent trained in fp16 and then cumulatively add our proposed methods one-by-one. The performance, averaged over environments and seeds, is shown in Figure 3. As we can see, all proposed techniques improve the performance individually, suggesting that they all are needed to achieve the final result. In appendix E we further experiment with removing individual components from the final agent which uses all other modifications and again see that our methods contribute individually to the final performance.




Figure 3. Ablation experiment where we add our components one-by-one until the complete agent is reached. Performance is averaged over seeds and tasks. The proposed methods improve performance individually, suggesting that they all contribute to the final performance. Fp16 crashes and yields no rewards.




Figure 4. Simulation of training in other number formats. The x-axis denotes the number of bits in the significand. Low-precision simulation is performed with qtorch (Zhang et al., 2019). Results are averaged across tasks. Performance decreases monotonically with fewer bits, first gracefully but then dramatically.

4.5. Other Numerical Formats

In addition to performing experiments in fp16, we also consider training in other numerical formats via the qtorch simulator (Zhang et al., 2019). Qtorch allows one to simulate low-precision training by quantizing tensors to some predefined format between calls to the PyTorch backend. Specifically, we start with the fp16 representation and vary the number of bits in the significand while retaining 5 bits for the exponent. For each resulting numerical format, we train SAC with our modifications. Results are shown in Figure 4. We see that the performance degrades with decreased numerical precision, first gracefully but then dramatically when using 5 bits in the significand.




Figure 5. Learning curves for SAC trained from pixels using standard fp32 training and fp16 training with our modifications. Average performance is close, demonstrating that low-precision RL is feasible from raw images too.

4.6. RL from Pixels

Until now we have focused on the traditional RL setting of learning from the states of the tasks (i.e. proprioceptive information), which is the original setup for SAC (Haarnoja et al., 2018). It is also possible to perform RL from raw pixel observations (Laskin et al., 2020b; Kostrikov et al., 2020). This setup is more compute-intensive, making low-precision methods very attractive. To this end, we replace the states of our original setup with raw pixel observations processed by a small convolutional network and perform end-to-end RL from pixels. For the convolutional encoder, we follow the general network design of (Kostrikov et al., 2020; Yarats et al., 2019; Laskin et al.) – a few convolutional layers (we use four as (Kostrikov et al., 2020)) followed by a linear layer which feeds into a 50-dimensional layer-normalization module whose output is used instead of the environment states. Occasionally we have found the internal variance calculations of layer-normalization to overflow in fp16, to remedy this, we add a weight standardization (Qiao et al., 2019) for the linear layer and downscale its output. This does not change the semantics as layer-normalization is invariant under rescaling. The issue could possibly be addressed at the GPU kernel level, but since the cudnn kernels are not open source we defer such studies to future work. Images are resized to 84-by-84 tensors and we follow (Kostrikov et al., 2020) regarding frame-stacking, image augmentation and hyperparameters (importantly learning rate is increased to $1e-3$), see Appendix G for details.

In Figure 5, we compare the performance between training

from pixels in fp32 and fp16, the latter using our modifications from section 3. Again, we see that the performance is very close, despite one network using half-precision numbers. This demonstrates the feasibility of RL from pixels in low precision too, where high-dimensional image input can result in compute-intensive workloads.

4.7. Performance

We now measure the performance improvements obtained from low-precision training, varying the batch size and the width (i.e. number of filters) of the convolutional layers to simulate different computational demands. We use all our proposed methods for low-precision training, larger improvements can likely be obtained for environments that work well without all methods. We here focus on RL from pixels, Appendix H contains the measurements for learning from states. When training the agent, there are three main components of the compute time: 1) simulating the environment, 2) sampling batches and 3) GPU computation. To achieve optimal performance, these operations should be asynchronous; however, the codebase of (Yarats & Kostrikov, 2020) is synchronous. As asynchronous learning is well-understood and not the focus of our work, we only measure the compute time of the third component, using one fixed batch and performing multiple training steps. The time is averaged over 500 gradient updates (and occasional momentum updates for the target network) with a warm started Tesla V100 GPU. We perform measurements with PyTorch tools, see Appendix H for tooling details.

Table 2. Time (milliseconds) for processing one minibatch as a function of network width and batch size (bsize). As computational demands increase the improvement from fp16 training exceeds a factor of two. Measurements are performed on the Cheetah task.

width / bsize	32/512	32/1024	64/512	64/1024
fp32	92.98	181.53	188.96	373.43
fp16 (ours)	76.46	127.10	93.76	171.12
improvement	1.22	1.43	2.02	2.18

Table 3. Memory (GB) consumed as a function of network width and batch size (bsize). Across all batch sizes and network widths savings are close to a factor of two. Measurements are performed on the Cheetah task.

bsize/ # filters	512/32	1024/32	512/64	1024/64
fp32	2.55	4.94	4.23	8.21
fp16 (ours)	1.36	2.61	2.28	4.37
improvement	1.87	1.89	1.86	1.88

In Table 2, we compare the time per update for different network widths and batch size. We see that the speedup increases to more than 2x as the computational demands increase. The reason that performance changes is likely that the smallest configuration does not saturate available compute resources of the Tesla V100, however factors such as memory hierarchy and pipeline balance might also play a role. In Table 3, we show the memory improvements of training in fp32 vs. fp16. The memory improvements stay around 1.87x for all network widths, demonstrating dramatic improvements that can enable larger models on edge devices and the small memory overhead of our methods.

4.8. Investigating Gradient Scales

The good performance of our proposed methods indicates that the numerical stability of policy parametrization, precision in gradient accumulation, and the dynamic range of the gradients are important issues in low-precision RL. We here seek to illustrate the wide gradient distribution directly. In Figure 6, we show the distribution of gradients for the cheetah task (trained in fp32) after 250,000 environment steps. Note that both axes use a logarithmic scale. The gradients have many orders of magnitude of dynamic scale, and when Adam computes the square of these even more dynamic range is needed. This suggests that gradient scale is an important issue in low precision RL.

4.9. Further Experiments

We have demonstrated that our methods can match fp32 performance for the tuned default SAC parameters of Yarats & Kostrikov (2020). As we did not tune these parameters ourselves, controlling for effects of hyperparameter tuning




Figure 6. Histogram over the gradients for the critic and actor networks. Note that both the x-axis and the y-axis are logarithmic. The gradients have a dynamic range spanning orders of magnitude, which can be problematic to represent in low precision.

is not straightforward (Jordan et al., 2020). To demonstrate that our methods are stable for different hyperparameters, we consider experiments with random hyperparameters in Appendix F. These show that our method still matches fp32 rewards for such random parameters, thus demonstrating the parameter stability of our methods.

5. Related Work

Accelerating RL. There are many proposed approaches for accelerating RL. (Mnih et al., 2016; Babaeizadeh et al., 2016) demonstrate the efficacy of asynchronous training, whereas distributed training has received more attention lately (Horgan et al., 2018; Espeholt et al., 2018; Wijmans et al., 2019; Espeholt et al., 2019; Kapturowski et al., 2018). Distributed RL training has received support from the Rllib project (Moritz et al., 2018; Liang et al., 2017), other notable codebases are Rlpyt (Stooke & Abbeel, 2019) and torchbeast (Küttler et al., 2019). Additionally, (Petrenko et al., 2020) and (Liu et al., 2020a) have recently demonstrated algorithms for high throughput RL on a single compute node. (Dalton et al., 2019) has investigated simulating the environment on the GPU for acceleration. (Lau et al., 2002) proposed vector quantization for traditional tabular reinforcement learning – i.e., with no deep learning component. The only work on accelerating deep RL with low precision known to us is the pre-print of (Krishnan et al., 2019). It differs from and/or complements our work by 1) studying post-training quantization, which is not the focus of our work, 2) using mixed precision, which we experimentally demonstrate does not work in general.

Low-Precision Deep Learning. Low-precision deep learning goes back at least to (Gupta et al., 2015), which has inspired large amounts of subsequent work. Among the most noteworthy approaches is mixed precision (Micikevicius et al., 2017), where full and half precision are used together. Such settings often necessitate loss scaling to avoid the gradients underflowing in the backward pass. Furthermore, it requires a larger set of instructions at the hardware level and also moving parameters between e.g. fp32 and

fp16 representations. Other relevant approaches include binary weights (Courbariaux et al., 2015), fixed point (Lin et al., 2016), B-float (Kalamkar et al., 2019), block floating point (Song et al., 2017). 8-bit (Wang et al., 2018), low-precision accumulation (Ni et al., 2020) or integer arithmetic (Jacob et al., 2018; Das et al., 2018). Another strategy is to average the low-precision weights into full precision (Yang et al., 2019) or to perform stochastic rounding (Hopkins et al., 2020). For recent theoretical work on quantization, see (Chen et al., 2020). Kahan summation has recently been proposed in deep learning for gradient accumulation (Zamirai et al., 2020; Han et al., 2019), but not for momentum updates as we do. Post-training quantization is also a popular strategy (Fang et al., 2020), but not the focus here.

6. Conclusions

In this paper, we demonstrate the feasibility and potential performance gains of low-precision RL. Whereas naïve strategies from supervised learning fail we propose six modifications to the SAC agent (Haarnoja et al., 2018) which improve numerical stability. With our modifications, fp16 training matches full precision rewards while improving memory footprint and compute time. Future directions include mixed precision and extending our methods to more agents. We hope that our methods and code provide a starting point for researchers interested in low-precision RL.

Acknowledgements

This research is supported in part by the grants from the National Science Foundation (III-1618134, III-1526012, IIS1149882, IIS-1724282, and TRIPODS- 1740822), the Office of Naval Research DOD (N00014- 17-1-2175), Bill and Melinda Gates Foundation. We are thankful for generous support by Zillow and SAP America Inc. This material is based upon work supported by the National Science Foundation under Grant Number CCF-1522054. We are also grateful for support through grant AFOSR-MURI (FA9550-18-1-0136). This work was supported by a gift from SambaNova Systems, Inc. Any opinions, findings, conclusions, or recommendations expressed here are those of the authors and do not necessarily reflect the views of the sponsors. We thank Rich Bernstein, Yiwei Bai, Wenting Zhao and Ziwei Liu for help with the manuscript. We also thank David Eriksson for pointers regarding numerics.

References

Ba, J. L., Kiros, J. R., and Hinton, G. E. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.

Babaeizadeh, M., Frosio, I., Tyree, S., Clemons, J., and Kautz, J. Reinforcement learning through asyn-

chronous advantage actor-critic on a gpu. *arXiv preprint arXiv:1611.06256*, 2016.

Bellemare, M. G., Candido, S., Castro, P. S., Gong, J., Machado, M. C., Moitra, S., Ponda, S. S., and Wang, Z. Autonomous navigation of stratospheric balloons using reinforcement learning. *Nature*, 588(7836):77–82, 2020.

Berner, C., Brockman, G., Chan, B., Cheung, V., Debiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C., et al. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*, 2019.

Carilli, M. Automatic mixed precision package - torch.cuda.amp. <https://pytorch.org/docs/stable/amp.html>, 2020.

Chen, J., Gai, Y., Yao, Z., Mahoney, M. W., and Gonzalez, J. E. A statistical framework for low-bitwidth training of deep neural networks. *arXiv preprint arXiv:2010.14298*, 2020.

Courbariaux, M., Bengio, Y., and David, J.-P. Binaryconnect: Training deep neural networks with binary weights during propagations. *Advances in neural information processing systems*, 28:3123–3131, 2015.

Dalton, S., Garland, M., and Frosio, I. Accelerating reinforcement learning through gpu atari emulation. 2019.

Das, D., Mellempudi, N., Mudigere, D., Kalamkar, D., Avancha, S., Banerjee, K., Sridharan, S., Vaidyanathan, K., Kaul, B., Georganas, E., et al. Mixed precision training of convolutional neural networks using integer operations. *arXiv preprint arXiv:1802.00930*, 2018.

De Sa, C., Feldman, M., Ré, C., and Olukotun, K. Understanding and optimizing asynchronous low-precision stochastic gradient descent. In *Proceedings of the 44th Annual International Symposium on Computer Architecture*, pp. 561–574, 2017.

Engstrom, L., Ilyas, A., Santurkar, S., Tsipras, D., Janoos, F., Rudolph, L., and Madry, A. Implementation matters in deep policy gradients: A case study on ppo and trpo. *arXiv preprint arXiv:2005.12729*, 2020.

Espeholt, L., Soyer, H., Munos, R., Simonyan, K., Mnih, V., Ward, T., Doron, Y., Firoiu, V., Harley, T., Dunning, I., et al. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. In *International Conference on Machine Learning*, pp. 1407–1416. PMLR, 2018.

Espeholt, L., Marinier, R., Stanczyk, P., Wang, K., and Michalski, M. Seed rl: Scalable and efficient deep-rl with accelerated central inference. *arXiv preprint arXiv:1910.06591*, 2019.

- Fang, J., Shafiee, A., Abdel-Aziz, H., Thorsley, D., Georgiadis, G., and Hassoun, J. H. Post-training piecewise linear quantization for deep neural networks. In *European Conference on Computer Vision*, pp. 69–86. Springer, 2020.
- Gong, J., Shen, H., Zhang, G., Liu, X., Li, S., Jin, G., Maheshwari, N., Fomenko, E., and Segal, E. Highly efficient 8-bit low precision inference of convolutional neural networks with intelcaffe. In *Proceedings of the 1st on Reproducible Quality-Efficient Systems Tournament on Co-designing Pareto-efficient Deep Learning*, pp. 1. 2018.
- Gupta, S., Agrawal, A., Gopalakrishnan, K., and Narayanan, P. Deep learning with limited numerical precision. In *International Conference on Machine Learning*, pp. 1737–1746, 2015.
- Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 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.
- Han, R., Si, M., Demmel, J., and You, Y. Auto-precision scaling for distributed deep learning. *arXiv preprint arXiv:1911.08907*, 2019.
- Henderson, P., Islam, R., Bachman, P., Pineau, J., Precup, D., and Meger, D. Deep reinforcement learning that matters. *arXiv preprint arXiv:1709.06560*, 2017.
- Hopkins, M., Mikaitis, M., Lester, D. R., and Furber, S. Stochastic rounding and reduced-precision fixed-point arithmetic for solving neural ordinary differential equations. *Philosophical Transactions of the Royal Society A*, 378(2166):20190052, 2020.
- Horgan, D., Quan, J., Budden, D., Barth-Maron, G., Hessel, M., Van Hasselt, H., and Silver, D. Distributed prioritized experience replay. *arXiv preprint arXiv:1803.00933*, 2018.
- Jacob, B., Kligys, S., Chen, B., Zhu, M., Tang, M., Howard, A., Adam, H., and Kalenichenko, D. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2704–2713, 2018.
- Jaderberg, M., Czarnnecki, W. M., Dunning, I., Marris, L., Lever, G., Castaneda, A. G., Beattie, C., Rabinowitz, N. C., Morcos, A. S., Ruderman, A., et al. Human-level performance in 3d multiplayer games with population-based reinforcement learning. *Science*, 364(6443):859–865, 2019.
- Jordan, S., Chandak, Y., Cohen, D., Zhang, M., and Thomas, P. Evaluating the performance of reinforcement learning algorithms. In *International Conference on Machine Learning*, pp. 4962–4973. PMLR, 2020.
- Kahan, W. Further remarks on reducing truncation errors, *commun. Assoc. Comput. Mach*, 8:40, 1965.
- Kahan, W. Ieee standard 754 for binary floating-point arithmetic. *Lecture Notes on the Status of IEEE*, 754(94720-1776):11, 1996.
- Kalamkar, D., Mudigere, D., Mellempudi, N., Das, D., Banerjee, K., Avancha, S., Vooturi, D. T., Jammalamadaka, N., Huang, J., Yuen, H., et al. A study of bfloat16 for deep learning training. *arXiv preprint arXiv:1905.12322*, 2019.
- Kapturowski, S., Ostrovski, G., Quan, J., Munos, R., and Dabney, W. Recurrent experience replay in distributed reinforcement learning. In *International conference on learning representations*, 2018.
- Kendall, A., Hawke, J., Janz, D., Mazur, P., Reda, D., Allen, J.-M., Lam, V.-D., Bewley, A., and Shah, A. Learning to drive in a day. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 8248–8254. IEEE, 2019.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kostrikov, I., Yarats, D., and Fergus, R. Image augmentation is all you need: Regularizing deep reinforcement learning from pixels. 2020.
- Krishnan, S., Chitlangia, S., Lam, M., Wan, Z., Faust, A., and Reddi, V. J. Quantized reinforcement learning (quarl). *arXiv preprint arXiv:1910.01055*, 2019.
- Küttler, H., Nardelli, N., Lavril, T., Selvatici, M., Sivakumar, V., Rocktäschel, T., and Grefenstette, E. Torchbeast: A pytorch platform for distributed rl. *arXiv preprint arXiv:1910.03552*, 2019.
- Laskin, M., Lee, K., Stooke, A., Pinto, L., Abbeel, P., and Srinivas, A. Reinforcement learning with augmented data. *arXiv:2004.14990*.
- Laskin, M., Lee, K., Stooke, A., Pinto, L., Abbeel, P., and Srinivas, A. Reinforcement learning with augmented data. *arXiv preprint arXiv:2004.14990*, 2020a.

- Laskin, M., Srinivas, A., and Abbeel, P. Curl: Contrastive unsupervised representations for reinforcement learning. In *Proceedings of the 37th Annual International Conference on Machine Learning (ICML)*, 2020b.
- Lau, H., Mak, K., and Lee, I. Adaptive vector quantization for reinforcement learning. *IFAC Proceedings Volumes*, 35(1):493–498, 2002.
- Lee, A. X., Nagabandi, A., Abbeel, P., and Levine, S. Stochastic latent actor-critic: Deep reinforcement learning with a latent variable model. *arXiv preprint arXiv:1907.00953*, 2019.
- Liang, E., Liaw, R., Moritz, P., Nishihara, R., Fox, R., Goldberg, K., Gonzalez, J. E., Jordan, M. I., and Stoica, I. Rllib: Abstractions for distributed reinforcement learning. arxiv e-prints, art. *arXiv preprint arXiv:1712.09381*, 2017.
- Lin, D., Talathi, S., and Annapureddy, S. Fixed point quantization of deep convolutional networks. In *International conference on machine learning*, pp. 2849–2858, 2016.
- Liu, I.-J., Yeh, R., and Schwing, A. High-throughput synchronous deep rl. *Advances in Neural Information Processing Systems*, 33, 2020a.
- Liu, Y., Jiang, Z., Zhang, S., and Xu, S. Deep reinforcement learning-based beam tracking for low-latency services in vehicular networks. In *ICC 2020-2020 IEEE International Conference on Communications (ICC)*, pp. 1–7. IEEE, 2020b.
- Micikevicius, P., Narang, S., Alben, J., Diamos, G., Elsen, E., Garcia, D., Ginsburg, B., Houston, M., Kuchaiev, O., Venkatesh, G., et al. Mixed precision training. *arXiv preprint arXiv:1710.03740*, 2017.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., and Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pp. 1928–1937, 2016.
- Moritz, P., Nishihara, R., Wang, S., Tumanov, A., Liaw, R., Liang, E., Elibol, M., Yang, Z., Paul, W., Jordan, M. I., et al. Ray: A distributed framework for emerging {AI} applications. In *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*, pp. 561–577, 2018.
- Ni, R., Chu, H.-m., Castañeda, O., Chiang, P.-y., Studer, C., and Goldstein, T. Wrapnet: Neural net inference with ultra-low-resolution arithmetic. *arXiv preprint arXiv:2007.13242*, 2020.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, pp. 8026–8037, 2019.
- Petrenko, A., Huang, Z., Kumar, T., Sukhatme, G., and Koltun, V. Sample factory: Egocentric 3d control from pixels at 100000 fps with asynchronous reinforcement learning. In *International Conference on Machine Learning*, pp. 7652–7662. PMLR, 2020.
- Qiao, S., Wang, H., Liu, C., Shen, W., and Yuille, A. Weight standardization. *arXiv preprint arXiv:1903.10520*, 2019.
- Reuther, A., Michaleas, P., Jones, M., Gadepally, V., Samsi, S., and Kepner, J. Survey and benchmarking of machine learning accelerators. *arXiv preprint arXiv:1908.11348*, 2019.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- Song, Z., Liu, Z., and Wang, D. Computation error analysis of block floating point arithmetic oriented convolution neural network accelerator design. *arXiv preprint arXiv:1709.07776*, 2017.
- Stooke, A. and Abbeel, P. rlpyt: A research code base for deep reinforcement learning in pytorch. *arXiv preprint arXiv:1909.01500*, 2019.
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. MIT press, 2018.
- 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.
- Tomy, M., Lacerda, B., Hawes, N., and Wyatt, J. L. Battery charge scheduling in long-life autonomous mobile robots. In *2019 European Conference on Mobile Robots (ECMR)*, pp. 1–6. IEEE, 2019.
- Wang, N., Choi, J., Brand, D., Chen, C.-Y., and Gopalakrishnan, K. Training deep neural networks with 8-bit floating point numbers. In *Advances in neural information processing systems*, pp. 7675–7684, 2018.
- Watkins, C. J. C. H. Learning from delayed rewards. 1989.
- Whitehead, N. and Fit-Florea, A. Precision & performance: Floating point and ieee 754 compliance for nvidia gpus. *rn (A+ B)*, 21(1):18749–19424, 2011.

- Wijmans, E., Kadian, A., Morcos, A., Lee, S., Essa, I., Parikh, D., Savva, M., and Batra, D. Dd-ppo: Learning near-perfect pointgoal navigators from 2.5 billion frames. *arXiv preprint arXiv:1911.00357*, 2019.
- Yang, G., Zhang, T., Kirichenko, P., Bai, J., Wilson, A. G., and De Sa, C. Swalp: Stochastic weight averaging in low precision training. In *International Conference on Machine Learning*, pp. 7015–7024. PMLR, 2019.
- Yarats, D. and Kostrikov, I. Soft actor-critic (sac) implementation in pytorch. https://github.com/denisyarats/pytorch_sac, 2020.
- Yarats, D., Zhang, A., Kostrikov, I., Amos, B., Pineau, J., and Fergus, R. Improving sample efficiency in model-free reinforcement learning from images. 2019.
- Zamirai, P., Zhang, J., Aberger, C. R., and De Sa, C. Revisiting bfloat16 training. *arXiv preprint arXiv:2010.06192*, 2020.
- Zhan, A., Zhao, P., Pinto, L., Abbeel, P., and Laskin, M. A framework for efficient robotic manipulation. *arXiv preprint arXiv:2012.07975*, 2020.
- Zhang, T., Lin, Z., Yang, G., and Sa, C. D. Qpytorch: A low-precision arithmetic simulation framework, 2019.
- Zhu, H., Yu, J., Gupta, A., Shah, D., Hartikainen, K., Singh, A., Kumar, V., and Levine, S. The ingredients of real-world robotic reinforcement learning. *arXiv preprint arXiv:2004.12570*, 2020.