

Learning optimal policies through contact in differentiable simulation

Anonymous authors

Paper under double-blind review

Abstract

Model-Free Reinforcement Learning (MFRL) has garnered significant attention for its effectiveness in continuous motor control tasks. However, its limitations become apparent in high-dimensional problems, often leading to suboptimal policies even with extensive training data. Conversely, First-Order Model-Based Reinforcement Learning (FO-MBRL) methods harnessing differentiable simulation offer more accurate gradients but are plagued by instability due to exploding gradients arising from the contact approximation model. We propose Adaptive Horizon Actor Critic (AHAC), a massively parallel FO-MBRL approach that truncates trajectory gradients upon encountering stiff contact, resulting in more stable and accurate gradients. We experimentally show this on a variety of simulated locomotion tasks, where our method achieves up to 91 % higher asymptotic episodic reward than state-of-the-art MFRL algorithms while also exhibiting lower variance and less hyperparameter sensitivity than prior FO-MBRL methods. Moreover, our method scales to high-dimensional motor control tasks while maintaining better wall-clock-time efficiency. We believe the ability to learn high-performance policies in a few minutes enables new opportunities to scale RL for robot motor control. <https://adaptive-horizon-actor-critic.github.io/>

1 Introduction

Reinforcement Learning (RL) has achieved remarkable success in complex tasks, such as Atari games (Mnih et al., 2013), Minecraft (Hafner et al., 2023) and Go (Silver et al., 2017). Combined with the Policy Gradients Theorem (Sutton et al., 1999), we can derive approaches for solving continuous motor control tasks. Although some of these Model-Free Reinforcement Learning (MFRL) approaches have achieved impressive results (Hwangbo et al., 2017; Akkaya et al., 2019; Hwangbo et al., 2019), they suffer from subpar sample efficiency, limiting their practical utility (Amos et al., 2021). Consequently, addressing this limitation has been a central research focus for the past few years.

An alternative approach, Model-Based Reinforcement Learning (MBRL), seeks to model the environment’s dynamics for improved efficiency. However, most MBRL methods still rely on experience data to learn dynamics and often produce suboptimal policies compared to MFRL. To enhance sample efficiency, high-performance physics simulators can be employed, offering abundant training data in ideal scenarios. When combined with computationally efficient MFRL methods, these simulators enable quick training of robots for tasks like walking (Rudin et al., 2022). However, a question remains: even with extensive data, can MFRL effectively tackle high-dimensional motor control problems?

Despite the success of MFRL methods in simulation, real-world applications face the challenge of the sim-to-real gap. To bridge this gap, one approach involves adapting the simulation model continuously with real-world data to achieve optimal task performance (Kumar et al., 2021). This is conceptually similar to learning a model, as in MBRL, making it tempting to learn the policy using first-order methods. This has been explored in the domain of model-based control literature, where the model is used for trajectory generation for applications such as autonomous driving (Kabzan et al., 2019) and agile quadrotor maneuvers (Kaufmann et al., 2020). However, using first-order methods to learn feedback policies in standard RL settings has received limited attention. Non-differentiable contact point discontinuities in available simulation models




Figure 1: We find that FO-MBRL methods suffer from high dynamics gradients $\|\nabla f(s, a)\| \gg 0$ which often arise from stiff contact approximation. Our proposed method, AHAC, truncates model-based trajectories at the point of and during stiff contact, thus avoiding both the gradient bias and learning instability exhibited by previous methods using differentiable simulation.

have been a major hurdle, leading to the development of differentiable simulators (Hu et al., 2019a; Freeman et al., 2021; Heiden et al., 2021; Xu et al., 2021).

These simulators have enabled a new family of FO-MBRL algorithms that can efficiently learn complex control tasks. Short Horizon Actor Critic (SHAC) (Xu et al., 2022) is such an approach that utilises the popular actor-critic paradigm (Konda & Tsitsiklis, 1999). The actor is trained in a first-order fashion, while the critic is trained model-free. This allows SHAC to learn through the highly non-convex landscape by using the critic as a smooth surrogate of the cumulative reward objective and avoiding exploding gradients by employing short-horizon rollouts. While SHAC boasts incredible sample efficiency when compared against MFRL, it is also brittle, exhibits higher learning instability, and requires extensive hyper-parameter tuning (Suh et al., 2022).

In this study, we attempt to address those issues and shift our focus from sample efficiency to the asymptotic performance of FO-MBRL methods in massively parallel differentiable simulations. We aim to answer the following questions:

1. What causes learning instability in FO-MBRL approaches such as SHAC? Our analysis reveals that first-order methods exhibit high empirical bias when estimating gradients through sample-limited Monte-Carlo approximation, hindering efficiency and resulting in suboptimal policies. This bias is primarily driven by the high magnitude dynamical gradients ($\|\nabla f(s, a)\| \gg 0$) arising from stiff contact approximation.
2. Can FO-MBRL methods outperform MFRL in finding optimal policies? We introduce Adaptive Horizon Actor Critic (AHAC), a first-order model-based algorithm that mitigates gradient issues during stiff contact by adapting its trajectory rollout horizon (Figure 1). Experimentally, we show that it is capable of achieving up to 91% more asymptotic reward in comparison to model-free approaches across complex locomotion tasks.
3. Which methods are suitable for scaling to high-dimensional motor control tasks? We find that AHAC exhibits lower variance during training, offering stability and gradient accuracy, allowing it to scale effectively to high-dimensional motor control tasks with action dimension $\mathcal{A} = \mathbb{R}^{152}$.

2 Preliminaries

In this paper, we study discrete-time, finite-horizon, fully-observable reinforcement learning problems where the state of the system is defined as $s \in \mathbb{R}^n$, actions are defined as $a \in \mathbb{R}^m$, and the dynamics are governed by the function $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$. Unlike traditional RL formulations, here we assume that the dynamics (i.e., transition function) are deterministic. At each timestep t , we sample an action from a stochastic policy

$\mathbf{a}_t \sim \pi_{\boldsymbol{\theta}}(\cdot|\mathbf{s}_t)$, which is parameterised by some parameters $\boldsymbol{\theta} \in \mathbb{R}^d$, and in turn, we receive a reward from $r : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$. We can define the H-step return as:

$$R_H(\mathbf{s}_1, \boldsymbol{\theta}) = \sum_{h=1}^H r(\mathbf{s}_h, \mathbf{a}_h) \quad s.t. \quad \mathbf{s}_{h+1} = f(\mathbf{s}_h, \mathbf{a}_h) \quad \mathbf{a}_h \sim \pi_{\boldsymbol{\theta}}(\cdot|\mathbf{s}_h)$$

As is typical in RL, the objective of the policy is to maximise the cumulative reward:

$$\max_{\boldsymbol{\theta}} J(\boldsymbol{\theta}) := \max_{\boldsymbol{\theta}} \mathbb{E}_{\substack{\mathbf{s}_1 \sim \rho \\ \mathbf{a}_h \sim \pi(\cdot|\mathbf{s}_h)}} [R_H(\mathbf{s}_1, \boldsymbol{\theta})] \quad (1)$$

where ρ is the initial state distribution. Without loss of generality, we make our work easier to follow by making the following assumption:

Assumption 2.1. We assume that ρ is a dirac-delta distribution.

Similar to prior work Duchi et al. (2012); Berahas et al. (2022); Suh et al. (2022), we are trying to exploit the smoothing properties of stochastic optimisation on the landscape of our optimisation objective. Following recent successful deep-learning approaches to MFRL (Schulman et al., 2017; Haarnoja et al., 2018), we assume:

Assumption 2.2. We assume that our policy is a parameterised diagonal Gaussian $\pi_{\boldsymbol{\theta}}(\cdot|\mathbf{s}) = \mathcal{N}(\mu_{\boldsymbol{\theta}}(\mathbf{s}), \Sigma_{\boldsymbol{\theta}}(\mathbf{s}))$ where $\mu_{\boldsymbol{\theta}} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ gives a multivariate mean and $\Sigma_{\boldsymbol{\theta}} : \mathbb{R}^n \rightarrow \mathbb{R}^{m \times m}$ gives a diagonal covariance matrix.

In order to solve our main optimisation problem in Equation 1, we consider using stochastic gradient estimates of $J(\boldsymbol{\theta})$. These can be obtained via zero-order and first-order methods. To guarantee the existence of $\nabla J(\boldsymbol{\theta})$, we need to make certain assumptions:

Definition 2.3. A function $g : \mathbb{R}^d \rightarrow \mathbb{R}^d$ has *polynomial growth* if there exists constants a, b such that $\forall \mathbf{z} \in \mathbb{R}^d, \|g(\mathbf{z})\| \leq a(1 + \|\mathbf{z}\|^b)$.

Assumption 2.4. To ensure gradients are well defined, we assume that the policy $\pi_{\boldsymbol{\theta}}(\cdot|\mathbf{s})$ is continuously differentiable $\forall \mathbf{s} \in \mathbb{R}^n, \forall \boldsymbol{\theta} \in \mathbb{R}^d$. Furthermore, the system dynamics f and reward r have polynomial growth.

2.1 Zeroth-Order Batch Gradient (ZOBG) estimates

These weak assumptions are sufficient to make $J(\boldsymbol{\theta})$ differentiable in expectation by simply taking samples of the function value in a zeroth-order fashion. This gives us estimates of $\nabla J(\boldsymbol{\theta})$ via the stochasticity introduced by π , as first shown in (Williams, 1992), and commonly referred to as the *Policy Gradient Theorem* (Sutton et al., 1999).

Definition 2.5. Given a sample of the H-step return $R_H(\mathbf{s}_1) = \sum_{h=1}^H r(\mathbf{s}_h, \mathbf{a}_h)$ following the policy π , we can estimate zero-order policy gradients via:

$$\nabla_{\boldsymbol{\theta}}^{[0]} J(\boldsymbol{\theta}) := \mathbb{E}_{\mathbf{a}_h \sim \pi_{\boldsymbol{\theta}}(\cdot|\mathbf{s}_h)} \left[R_H(\mathbf{s}_1) \sum_{h=1}^H \nabla_{\boldsymbol{\theta}} \log \pi_{\boldsymbol{\theta}}(\mathbf{a}_h|\mathbf{s}_h) \right] \quad (2)$$

Lemma 2.6. Under Assumptions 2.1 and 2.4, the ZOBG is an unbiased estimator of the stochastic objective $\mathbb{E}[\bar{\nabla}^{[0]} J(\boldsymbol{\theta})] = \nabla J(\boldsymbol{\theta})$ where $\bar{\nabla}^{[0]} J(\boldsymbol{\theta}) = \frac{1}{N} \sum_{n=1}^N \hat{\nabla}^{[0]} J(\boldsymbol{\theta})$ is the N -sample Monte-Carlo estimate of Eq. 2.

These zero-order policy gradients are known to have high variance, and one way to reduce their variance is by subtracting a baseline from the function estimates. Similar to (Suh et al., 2022), we do that by subtracting the return given by the noise-free policy rollout:

$$\nabla_{\boldsymbol{\theta}}^{[0]} J(\boldsymbol{\theta}) = \mathbb{E}_{\mathbf{a}_h \sim \pi_{\boldsymbol{\theta}}(\cdot|\mathbf{s}_h)} \left[\left(R_H(\mathbf{s}_1) - R_H^*(\mathbf{s}_1) \right) \sum_{h=1}^H \nabla_{\boldsymbol{\theta}} \log \pi_{\boldsymbol{\theta}}(\mathbf{a}_h|\mathbf{s}_h) \right] \quad R_H^*(\mathbf{s}_1) := \sum_{h=1}^H r(\mathbf{s}_h, \mu_{\boldsymbol{\theta}}(\mathbf{s}_h))$$

2.2 First-Order Batch Gradient (FOBG) estimates

Given access to a differentiable simulator (with contact approximation), one can directly compute the analytic gradients of $\nabla_{\theta} R_H(s_1)$ induced by the policy π :

$$\nabla_{\theta}^{[1]} J(\theta) := \mathbb{E}_{\mathbf{a}_h \sim \pi_{\theta}(\cdot | s_h)} [\nabla_{\theta} R_H(s_1)] \quad (3)$$

However, for these gradients to be well-defined, we need to make further assumptions:

Assumption 2.7. The system dynamics $f(\mathbf{s}, \mathbf{a})$ and the reward $r(\mathbf{s}, \mathbf{a})$ are continuously differentiable $\forall \mathbf{s} \in \mathbb{R}^n, \forall \mathbf{a} \in \mathbb{R}^m$.

3 Learning through contact

First-order gradients, as demonstrated in prior research, are asymptotically unbiased when $N \rightarrow \infty$ (Schulman et al., 2015). However, this ideal scenario is often impractical in real-world applications, leading to observed empirical bias, as indicated by (Suh et al., 2022). To illustrate this bias, we use the soft Heaviside function, a common tool for studying discontinuous functions in physics simulations:

$$R_1(s, \theta) = \bar{H}(a) = \begin{cases} 1 & a > \nu/2 \\ 2a/\nu & |a| \leq \nu/2 \\ 0 & a < -\nu/2 \end{cases} \quad (4)$$

where $a \sim \pi_{\theta}(\cdot) = \mathcal{N}(0, 1)$. This function can be shown to have an expected value equivalent to the error function: $\mathbb{E}_{\pi}[\bar{H}(a)] = \text{erf}(\nu/2; \sigma^2)$, which has $\nabla \mathbb{E}_{\pi}[\bar{H}(a)] = \frac{1}{\sqrt{2\pi}\sigma}$. Using FOBG, we obtain $\nabla_{\theta} \bar{H}(a) = 0$ in samples where $|a| > \nu/2$, which occurs with probability at least $\frac{\nu}{\sqrt{2\pi}\sigma}$. Since in practice we are limited in sample size, this translates to empirical bias that is inversely proportional to sample size, as shown in Figure 3. Notably, when ν approaches zero, we achieve more realistic behaviour, but we also increase the likelihood of obtaining incorrect FOBG, thus amplifying bias in stochastic scenarios.

By analysing the empirical bias from the perspective of bias and variance, we derive a practical upper bound:

Lemma 3.1. For an H-step stochastic optimisation problem under Assumptions 2.7, which also has Lipschitz-smooth policies $\|\nabla \pi_{\theta}(\mathbf{s} | \mathbf{s})\| \leq B_{\pi}$ and Lipschitz-smooth and bounded rewards $r(\mathbf{s}, \mathbf{a}) \leq \|\nabla r(\mathbf{s}, \mathbf{a})\| \leq B_r$ $\forall \mathbf{s} \in \mathbb{R}^n; \mathbf{a} \in \mathbb{R}^m; \theta \in \mathbb{R}^d$, then zero-order estimates remain unbiased. However, first-order gradient exhibit bias which is bounded by:

$$\left\| \mathbb{E} \left[\nabla_{\theta}^{[1]} J(\theta) \right] - \mathbb{E} \left[\nabla_{\theta}^{[0]} J(\theta) \right] \right\| \leq H^4 B_r^2 B_{\pi}^2 \mathbb{E}_{\mathbf{a} \sim \pi} \left[\prod_{t=1}^H \|\nabla f(\mathbf{s}_t, \mathbf{a}_t)\|^2 \right] \quad (5)$$

The proof can be found in Appendix A

We can ensure that the rewards are designed to meet the condition $r(\mathbf{s}, \mathbf{a}) \leq \|\nabla r(\mathbf{s}, \mathbf{a})\| \leq B_r$. The assumption over the policy $\|\nabla \pi_{\theta}(\cdot | \mathbf{s})\| \leq B_{\pi}$ is less straightforward if we are using high-capacity models such as large neural networks, but as is common practice in deep learning, we can normalise gradients to improve stability. However, giving any bounds over the dynamics $\|\nabla f(\mathbf{s}_t, \mathbf{a}_t)\|^2$ is difficult, yet is the biggest contributor to the variance in Equation 5.




Figure 2: Heaviside Function (Eq 4).




Figure 3: **Empirical bias of ZOFG and FOBG** when approximating the gradients of the Heaviside function.




Figure 5: The toy problem where the ball is shot against a wall to reach the blue box.

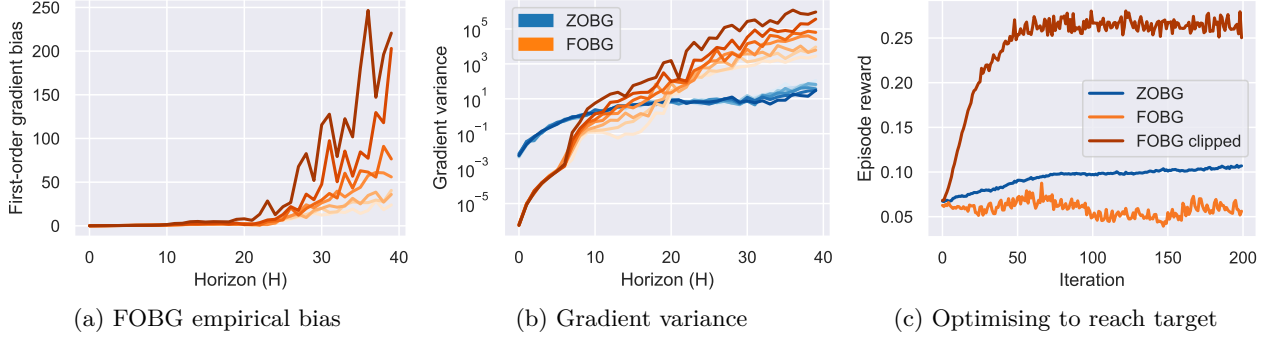


Figure 4: **Results from the toy ball problem.** Fig (a) shows the empirical bias of FOBG measures by comparing it to ZOBG. Fig (b) shows how the variance of the two gradient types evolves over time. Both Figs (a) and (b) use different shades to show different stiffness configurations of the simulation with darker shades designating stiffer (and more realistic) simulation. Fig (c) shows attempts at optimising the initial angle of the ball to hit the target.

To investigate the implications of Lemma 3.1, we constructed a simple scenario involving a ball rebounding off a wall and aiming to reach a target location, as illustrated in Figure 5. In this setup, the initial position $\mathbf{s}_1 = [x_1, y_1]$ and velocity of the ball are fixed. The objective is for the policy to learn the optimal initial orientation ψ following an additive Gaussian policy $\pi(\theta) = \theta + w$ where $w \sim \mathcal{N}(0, \sigma)$. With this, zero-order (policy) gradients from Equation 2 can be expressed as:

$$\nabla_{\theta}^{[0]} J(\theta) = \mathbb{E}_{\mathbf{a}_h \sim \pi(\theta)} [(R_H(\mathbf{s}_1) - R_H^*(\mathbf{s}_1)) \nabla \pi(\theta)] \approx \frac{1}{N} \frac{1}{\sigma^2} (R_H(\mathbf{s}_1) - R_H^*(\mathbf{s}_1)) w$$

We collect $N = 1024$ samples of each gradient type for each timestep with $H = 40$, starting from a randomly sampled angle $\psi \sim U(-\pi, \pi)$ for each environment. Figure 4a shows how the empirical bias of FOBG grows as H increases, validating the proposed lemma. The bias remains low until the ball encounters contact, at which point it starts growing exponentially. We also examined the variance of the gradients in Figure 4b, observing that ZOBG follow $\text{Var}[\nabla^{[0]} J(\theta)] \leq \frac{HB_r^2 B_\pi^2}{\sigma^2}$ first proposed by (Suh et al., 2022) - most importantly, they scale linearly with H . However, FOBG variance behaves similarly to the empirical bias bound described in Lemma 3.1. exhibiting lower variance at the beginning of the rollout but growing exponentially. This is due mostly to the stiff dynamics $\|\nabla f\| \gg 0$, which can be clearly seen at timestep $h = 6$ of Figure 4b where the ball first makes contact with the ground. Towards the end of the rollout, FOBG’s variance can be up to five orders of magnitude higher than ZOBG, which remains unaffected by stiff dynamics. Both the bias and variance issues become even more pronounced as the contact stiffness increases, indicated by the darker shades in the figures.

This high empirical bias and the resulting high variance have significant consequences for optimisation and policy learning. We attempted to optimize the toy problem using a reward that only considers the final state, where $r(\mathbf{s}_H, \mathbf{a}_H) = \|\mathbf{s}_H - \mathbf{s}_T\|_2^{-1}$, with \mathbf{s}_T representing the target state (with zero velocity). We find that when using FOBG in a high-stiffness scenario, the gradients have such high bias and variance that they do not make progress towards a solution, as shown in Figure 4c. In contrast, the unbiased ZOBG have lower variance and make progress towards a solution, albeit slowly.

This situation leads to a critical question: Is it possible to leverage the efficiency of FOBG for learning while still making progress toward a solution in the presence of high bias gradients, especially when the number of available samples (N) is limited? To address this challenge, we devised a solution inspired by the practice in deep learning, which involves clipping gradient norms. Specifically, we introduced a mechanism for clipping gradients of the dynamics at each time step, as described in the equations:

$$\begin{aligned} \tilde{\nabla}_{\mathbf{s}_h} f(\mathbf{s}_h, \mathbf{a}_h) &= gc(\nabla_{\mathbf{s}_h} f(\mathbf{s}_h, \mathbf{a}_h)) \quad \forall h \in [0, H] \\ \tilde{\nabla}_{\mathbf{a}_h} f(\mathbf{s}_h, \mathbf{a}_h) &= gc(\nabla_{\mathbf{a}_h} f(\mathbf{s}_h, \mathbf{a}_h)) \quad \forall h \in [0, H] \end{aligned} \quad gc(\mathbf{x}) := \begin{cases} \frac{\mathbf{x}}{\|\mathbf{x}\|_2} & \text{if } \|\mathbf{x}\|_2 > 1 \\ \mathbf{x} & \text{otherwise} \end{cases}$$

When we employed this approach in the optimisation process for the toy problem, using FOBG with the normalised gradients yielded a policy learning process that converged to a good solution much more quickly than ZOBG, as illustrated in Figure 4c. We believe this preserves most of the gradient information while reducing the magnitude and bias arising from the stiff dynamics, resulting in accurate and well-behaved gradients.

4 Adaptive Horizon Actor Critic (AHAC)

4.1 Learning through contact in a single environment

With a clearer understanding of stiff contact in differentiable simulators, we aim to develop a model-based algorithm employing FOBG for effective learning in infinite-horizon robotics tasks. Although we can apply the gradient clipping technique as above, in a multi-step problem, we can avoid the stiff gradients altogether. Consider a scenario where we have a 3-step trajectory ($H = 3$), and contact occurs at timestep $h = 3$, as shown in Figure 6. If we took the gradient normalisation approach, the gradients for $r(s_3, a_3)$ with a θ -parameterised policy where $a_h \sim \pi_\theta(\cdot|s_h)$ with respect to θ :

$$\begin{aligned} \nabla_\theta r(s_3, a_3) &= \nabla_{a_3} r(s_3, a_3) \nabla_\theta \pi_\theta(a_3|s_3) \\ &\quad + \nabla_{s_3} r(s_3, a_3) \tilde{\nabla}_{a_2} f(s_2, a_2) \nabla_\theta \pi_\theta(a_2|s_2) \\ &\quad + \nabla_{s_3} r(s_3, a_3) \tilde{\nabla}_{s_2} f(s_2, a_2) \tilde{\nabla}_{a_1} f(s_1, a_1) \nabla_\theta \pi_\theta(a_1|s_1) \end{aligned}$$

By definition of $gc(x)$ we will be losing gradient information. An alternative would be to cut gradients at $h = 2$. This would render $\nabla_\theta r(s_3) = 0$, preventing us from learning in a scenario such as our toy example in Section 3. However, in a multi-step decision-making problem, $r(s_3, a_3)$ still yields gradients:

$$\begin{aligned} \nabla_\theta \left[\sum_{h=1}^3 r(s_h, a_h) \right] &= \nabla_{a_3} r(s_3, a_3) \nabla_\theta \pi_\theta(a_3|s_3) \\ &\quad + \nabla_{a_2} r(s_2, a_2) \nabla_\theta \pi_\theta(a_2|s_2) + \nabla_{s_2} r(s_2, a_2) \nabla_{a_1} f(s_1, a_1) \nabla_\theta \pi_\theta(a_1|s_1) \\ &\quad + \nabla_{a_1} r(s_1, a_1) \nabla_\theta \pi_\theta(a_1|s_1) \end{aligned}$$

Notice that (1) the gradients of dynamics in contact do not appear and (2) how the gradients of timesteps in contact are akin to policy gradients. We call this technique *contact truncation* and view it as a blend between FOBG and ZOBG.

We present an FO-MBRL algorithm with an actor-critic architecture similar to SHAC (Xu et al., 2022). The critic, denoted as $V_\psi(s)$, is model-free and trained using TD(λ) over an H -step horizon from timestep t :

$$R_h(s_t) := \sum_{n=t}^{t+H-1} \gamma^{n-t} r(s_n, a_n) + \gamma^{t+H} V_\phi(s_{t+H}) \quad \hat{V}(s_t) := (1 - \lambda) \left[\sum_{h=1}^{H-t-1} \lambda^{h-1} R_h(s_t) \right] + \lambda^{H-t-1} R_H(s_t)$$

The critic loss becomes $\mathcal{L}(\psi)$, while the actor is trained using FOBG as in Equation 3, with the addition of the critic value estimate:

$$\mathcal{L}(\psi) := \sum_{h=t}^{t+H} \left\| V_\psi(s_h) - \hat{V}(s_h) \right\|_2^2 \quad (6) \quad J(\theta) := \sum_{h=t}^{t+H-1} \gamma^{h-t} r(s_h, a_h) + \gamma^t V_\psi(s_{t+T}) \quad (7)$$

Unlike fixed-horizon model-based rollouts in (Xu et al., 2022), our policy is rolled out until stiff contact is encountered, which can be determined in simulation. This results in a dynamic FO-MBRL algorithm that adjusts its horizon to avoid exploding gradients. However, not all contact results in high bias; therefore, we want to truncate only on stiff contact $\|\nabla f(s_t, a_t)\| > C$, where C is the contact stiffness parameter we set. We refer to this algorithm as Adaptive Horizon Actor Critic 1 (AHAC-1), which is designed for single environments but not suitable for vectorised environments (see Appendix C).

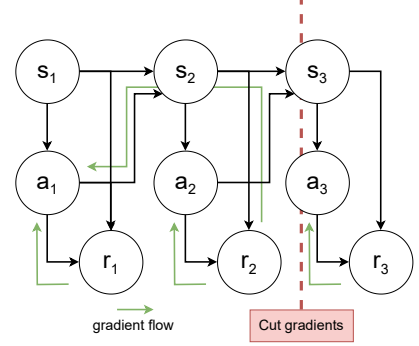


Figure 6: **H=3 trajectory with truncated gradients.** The green arrows indicate back-propagated gradients for AHAC-1.




Figure 7: **Comparison between SHAC and AHAC-1 ran on the Hopper task with only a single environment.** The left plot shows rewards achieved by the algorithms over five different random seeds, with the mean and variance plotted. The right plot is the moving window averaged mean horizon of both approaches. Note that even though SHAC has fixed horizons, they can still vary if the environment is terminated early due to early termination or episode end.

Using this approach, we can investigate if truncating gradients on contact yields better policy than naively cutting on fixed short-horizons. We compare SHAC and AHAC-1 on a simple contact-rich locomotion task, Hopper, a single-legged agent that obtains a reward for forward velocity (Figure 1). Both algorithms share the same hyperparameters, except for the ones related to horizons. SHAC uses a fixed $H = 32$, while AHAC-1 uses a maximum horizon of $H = 64$ with a contact threshold of $C = 500$. From Figure 7, we observe that AHAC-1 achieves a higher reward while exhibiting lower variance. Although difficult to analyse, we believe that our approach avoids local minima by adapting its horizon to avoid stiff gradients. On the other hand, SHAC gets pushed into local minima, which eventually results in policy collapse. Unfortunately, AHAC-1 cannot be applied to parallel vectorised environments due to the challenge of asynchronously truncating trajectories, leading to infinitely long compute graphs.

4.2 AHAC: A scalable approach for learning through contact

A straightforward solution to asynchronous truncation in AHAC-1 might involve adopting the short-horizon approach of SHAC and truncating the graph on stiff contact. Unfortunately, this did not yield any performance improvements. Instead, we investigate the impact of the horizon length on policy optimality. We find that contact-based tasks have an inherent optimal solution, for instance, in a locomotion task, in the form of an optimal gait pattern. We observe that an SHAC agent converges to a particular solution, often suboptimal, depending on the horizon parameter H . Importantly, we find that asymptotic performance is maximised when the horizon H matches the optimal (though unknown) gait frequency.

To empirically demonstrate this, we conducted an experiment by parameterizing SHAC with different H values for Ant locomotion tasks, where a quadruped robot seeks to maximise forward velocity rewards. As seen from the results in Figure 8, the gait period aligns with the horizon length until $H = 28$, after which it attempts to fit two gait patterns within a single H -timestep rollout. Moreover, we noticed that the asymptotic reward reaches its peak as the horizon-length H approaches what we believe to be the optimal gait period and displays the least variance across runs.




Figure 8: An ablation of short horizons H for the SHAC algorithm applies to Ant. Each run is trained until convergence for 5 random seeds.

From these observations, we glean two insights: (1) each task has an optimal model-based horizon length H that corresponds to the gait period, and (2) the associated optimal horizon results in the lowest variance between runs, supported by Lemma 3.1. We leverage these insights to generalise the AHAC-1 algorithm into a GPU-parallelisable approach, which we call AHAC. The critic training formulation remains the same as in Equation 6, but we introduce a new constrained objective for the actor:

$$J(\theta) := \sum_{h=t}^{t+H-1} \gamma^{h-t} r(\mathbf{s}_h, \mathbf{a}_h) + \gamma^t V_\psi(\mathbf{s}_{t+H}) \quad (8)$$

s.t. $\|\nabla f(\mathbf{s}_t, \mathbf{a}_t)\| \leq C \quad \forall t \in \{0, \dots, H\}$

In simple terms, this objective seeks to maximise the reward while ensuring that all contact stiffness remains below a predefined threshold. Building on the inspiration from AHAC-1, we can progressively increase the horizon as long as the constraint is satisfied. Using the Lagrangian formulation, we derive the dual problem:

$$\mathcal{L}(\theta, \phi) = \sum_{h=t}^{t+H-1} \gamma^{h-t} r(\mathbf{s}_h, \mathbf{a}_h) + \gamma^t V_\psi(\mathbf{s}_{t+H}) + \phi^T \left(\begin{bmatrix} \|\nabla f(\mathbf{s}_t, \mathbf{a}_t)\| \\ \vdots \\ \|\nabla f(\mathbf{s}_{t+H}, \mathbf{a}_{t+H})\| \end{bmatrix} - C \right) \quad (9)$$

By definition, $\phi_i = 0$ if the constraint is met and $\phi_i > 0$ otherwise. Thus, we can use ϕ to adapt the horizon, resulting in the full AHAC algorithm shown

in Algorithm 1. We train the critic until convergence is defined as a sufficiently small change in the last 5 critic training iterations, $\sum_{i=n-5}^n \mathcal{L}(\psi)$ where we take mini-batch samples from $(\mathbf{s}, \hat{V}(\mathbf{s})) \sim D$.

In practice, we find that truncating based simply on $\nabla f(\mathbf{s}_t, \mathbf{a}_t)$ is limiting since different tasks involve varying contact forces, and these forces often change throughout the learning process. Instead, we normalise the contact forces with modified acceleration per state dimension $\hat{\mathbf{a}}_t = \max(\mathbf{a}_t, 1)$, resulting in normalised contact forces $\hat{\nabla} f(\mathbf{s}_t, \mathbf{a}_t) = \frac{\nabla f(\mathbf{s}_t, \mathbf{a}_t)}{\hat{\mathbf{a}}_t}$. This allows us to use a single C parameter and still enables AHAC to adapt the horizon to different tasks. Additionally, as contact approximation forces are computed separately in differentiable simulators, we don't need to utilise the full Jacobian of the dynamics, as suggested by the equations above. Instead, we can use the Jacobian derived only from contact approximation. The differences between SHAC and AHAC are summarised in Appendix B.

5 Experiments

In this section, we aim to address the following key questions experimentally:

1. Does contact truncation enhance stability compared to prior differentiable simulator approaches?
2. Can first-order model-based (FO-MBRL) policies outperform zeroth-order model-free (ZO-MBRL) policies concerning asymptotic episodic reward?
3. Does AHAC remain sample and wall-clock time efficient similar to prior FO-MBRL algorithms?
4. Does AHAC scale to high-dimensional environments?

Algorithm 1: Adaptive Horizon Actor-Critic

Given: γ : discount rate

Given: α : learning rate

Given: C : contact threshold

Initialise learnable parameters $\theta, \psi, H, \phi = 0$

$t \leftarrow 0$

while *episode not done* **do**

/ rollout policy*

**/*

 Initialise buffer D

for $h = 0, 1, \dots, H$ **do**

$\mathbf{a}_{t+h} \sim \pi(\cdot | \mathbf{s}_{t+h}; \theta)$

$\mathbf{s}_{t+h+1} = f(\mathbf{s}_{t+h}, \mathbf{a}_{t+h})$

$D \leftarrow D \cup \{(\mathbf{s}_{t+h}, \mathbf{a}_{t+h}, \mathbf{r}_{t+h}, V_\psi(\mathbf{s}_{t+h+1}))\}$

/ train actor with Eq. 9*

**/*

$\theta \leftarrow \theta - \alpha \nabla_\theta \mathcal{L}(\theta, \phi)$

$\phi \leftarrow \phi - \alpha \nabla_\phi \mathcal{L}(\theta, \phi)$

$H \leftarrow H - \alpha \sum_{t=0}^H \phi_t$

/ train critic with Eq. 6*

**/*

while *not converged* **do**

 sample $(\mathbf{s}, \hat{V}(\mathbf{s})) \sim D$

$\psi \leftarrow \psi + \alpha \nabla_\psi \mathcal{L}(\psi)$

$t \leftarrow t + H$

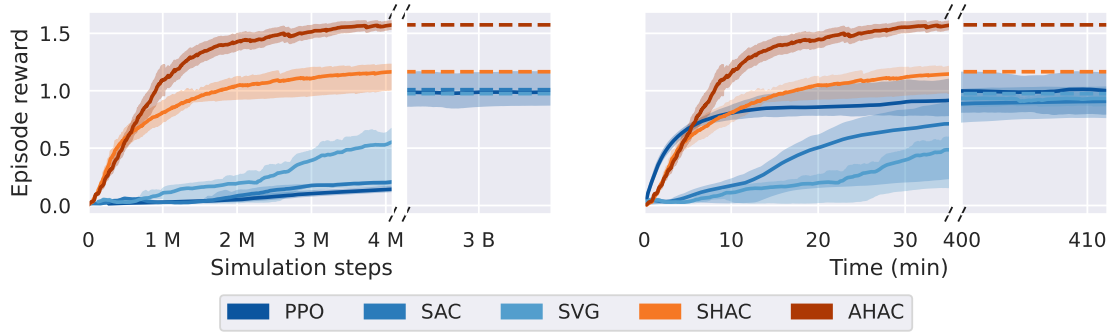


Figure 9: **Episodic rewards of the Ant task against both simulation steps and wall clock time.** The episodic reward is PPO-normalised, i.e., normalised by the asymptotic reward achieved by PPO at the end of training. The dashed lines represent the reward achieved by each respective algorithm at the end of their training runs.

Previous FO-MBRL algorithms utilising differentiable simulators, such as SHAC (Xu et al., 2022), face challenges related to instability arising from stiff contact. The previous section and Figure 7 specifically suggest that the single-environment version of AHAC reduces performance variance substantially. We now investigate whether these benefits persist when scaling AHAC to $N = 128$ parallel environments and applying it to a more complex task: Ant, a quadruped with symmetrical legs, $\mathcal{S} = \mathbb{R}^{37}$ and $\mathcal{A} = \mathbb{R}^8$. We compare AHAC to SHAC, its predecessor PPO, a state-of-the-art on-policy MFRL algorithm (Schulman et al., 2017), SAC, an off-policy MFRL algorithm (Haarnoja et al., 2018), and SVG, a FO-MBRL method that does not utilise a differentiable simulator but instead learns its model of the dynamics¹. A more explicit comparison of all of these approaches and more can be found in Section 6. We tune and train all algorithms on the Ant task until convergence with five different random seeds. Furthermore, we also train the MFRL baselines for 3B timesteps to investigate their asymptotic performance given practically infinite data.

Figure 9 provides insights into the asymptotic performance. It shows that AHAC maintains a lower performance variance compared to SHAC and achieves 57% higher reward than the closest model-free baseline. Remarkably, even when the MFRL algorithms, PPO and SAC, are allowed to train for an extended period (3 billion timesteps), they achieve asymptotically worse episodic reward when compared to the FO-MBRL approaches using differentiable simulation.

To answer questions (3) and (4) regarding efficiency and scalability, we conduct experiments across a variety of locomotion tasks. We again compare AHAC against SHAC, PPO, SAC, and SVG. Each algorithm has been tuned to perform well on all tasks using the same hyperparameter configuration, and all algorithms use the same parametrization for their actor and critic networks (full hyperparameters in Appendix D). We conduct five different seeded experiments per task on the same machine for fair comparison. The environments we consider are described below:

1. **Hopper**, a single-legged robot jumping only in one axis with $\mathcal{S} = \mathbb{R}^{11}$ and $\mathcal{A} = \mathbb{R}^3$.
2. **Anymal**, a more sophisticated quadruped with $\mathcal{S} = \mathbb{R}^{49}$ and $\mathcal{A} = \mathbb{R}^{12}$ modelled after the real robot (Hutter et al., 2016).
3. **Humanoid**, a classic contact-rich environment with $\mathcal{S} = \mathbb{R}^{76}$ and $\mathcal{A} = \mathbb{R}^{21}$ which requires extensive exploration to find a good policy.
4. **SNU Humanoid**, a version of Humanoid lower body where instead of joint torque control, the robot is controlled via $\mathcal{A} = \mathbb{R}^{152}$ muscles intended to challenge the scaling capabilities of algorithms.

¹To make performance comparable, we attempted to vectorize the original implementation of SVG and found that it did not scale well with an increased number of parallel environments. Therefore, the results presented in this paper are from the original single-environment version of the algorithm.



Figure 10: Locomotion environments (left to right): Hopper, Ant, Anymal, Humanoid and SNU Humanoid.




Figure 11: **Reward curves for all locomotion tasks against both simulation steps and training time.** The reward axis are normalised by the best reward achieved by PPO. Plotted are the means over 5 runs for each experiment, with the standard deviation shown in the shaded area. We further apply exponentially weighted smoothing with $\alpha = 0.98$ to increase legibility.

The main reward signal across all of these tasks is sustained forward-velocity with actions penalties. Environments are visualised in Figure 10.

	Hopper	Ant	Anymal	Humanoid	SNU Humanoid
PPO	1.00 ± 0.04	1.00 ± 0.14	1.00 ± 0.04	1.00 ± 0.06	1.00 ± 0.06
SAC	1.10 ± 0.02	1.01 ± 0.00	0.96 ± 0.05	0.98 ± 0.04	0.88 ± 0.06
SVG	0.95 ± 0.12	0.97 ± 0.11	0.79 ± 0.17	0.93 ± 0.16	0.66 ± 0.21
SHAC	1.13 ± 0.00	1.17 ± 0.11	1.26 ± 0.01	1.10 ± 0.07	1.59 ± 0.05
AHAC	1.21 ± 0.00	1.57 ± 0.05	1.43 ± 0.04	1.19 ± 0.03	1.91 ± 0.11

Table 1: Tabular results of the highest rewards achieved by each algorithm across all tasks. The results presented are mean and standard deviation across 5 random seeds. All algorithms have been trained until convergence with the exception of SVG which was only trained for 2M timesteps due to time limitations.

We compare the approaches against the number of simulation steps but also acknowledge that MFRL methods are computationally simpler and thus also provide results against wall-clock time. From the results in Figure 11, it can be clearly seen that AHAC maintains the sample efficiency of SHAC, significantly outperforming the other methods. On the simpler Hopper task, all algorithms achieve similar performance when compared

against wall-clock time. However, with more complex tasks, we observe that AHAC not only learns at a faster rate but is also capable of achieving higher asymptotic performance across all tasks. This gap only becomes larger as we turn to more high-dimensional environments, showcasing the remarkable scalability of our approach, where AHAC achieves 91% higher asymptotic reward than PPO on the SNU Humanoid task. Tabular results across all tasks are shown in Table 1, where most algorithms are trained until convergence.

6 Related work

In this section, we provide an overview of recent continuous control reinforcement learning (RL) methods, all of which follow the actor-critic paradigm (Konda & Tsitsiklis, 1999). The critic estimates the value of state-action pairs $Q(\mathbf{s}, \mathbf{a})$, and the actor learns optimal actions via $\max_{\mathbf{a}} Q(\mathbf{s}, \mathbf{a})$. While this section is intended as a review of related work, we also attempt to classify methods by their method of policy (actor) training, value estimator (critic) training, and their dynamics model $f(\mathbf{s}, \mathbf{a})$.

When no dynamics model is assumed, we are restricted to Model-Free Reinforcement Learning (MFRL) methods. We can take Monte-Carlo samples of the Policy Gradients Theorem to find $\nabla_{\theta} J(\theta)$ using Equation 2. This allows MFRL methods to learn a feedback policy that predicts the distribution of actions given the state. On-policy methods, like Proximal Policy Optimisation (PPO) (Schulman et al., 2017), learn using only the most recent samples following the policy. In contrast, off-policy approaches, such as Soft Actor-Critic (SAC) (Haarnoja et al., 2018), can use all previously collected data at the expense of memory requirements.

Alternatively, Model-Based Reinforcement Learning (MBRL) methods aim to leverage a model for learning. This model can be learned from experience data or assumed a priori. In a basic scenario, it serves as an additional source of return estimates for the critic, which can still be trained in a model-free manner (Janner et al., 2019). Alternatively, the model can be used to obtain simulated returns for the critic, which can be first-order back-propagated through, known as Model-based Value Expansion (MVE) (Feinberg et al., 2018). Actor training is more intricate in this context. It can be done using Policy Gradients augmented by model-generated data (Janner et al., 2019) or as part of a gradient-free planning actor (Hafner et al., 2019b). This family of approaches is termed Zeroth-Order MBRL (ZO-MBRL). Alternatively, the returns of trajectories can be used to backpropagate through the model (Hafner et al., 2019a; Byravan et al., 2020), and we refer to these methods as First-Order MBRL (FO-MBRL). Key recent work is summarised in Table 2.

	Algorithm	Policy Learning	Value Learning	Dynamics Model
MFRL	PPO (Schulman et al., 2017)	Zeroth-order	Model-free	-
	SAC (Haarnoja et al., 2018)	Zeroth-order	Model-free	-
ZO-MBRL	MBPO (Janner et al., 2019)	Zeroth-order	Model-free	Ensemble NN
	PlaNet (Hafner et al., 2019b)	Gradient-free	-	Probabilistic NN
FO-MBRL	MVE (Feinberg et al., 2018)	Zeroth-order	Model-based	Deterministic NN
	STEVE (Buckman et al., 2018)	Zeroth-order	Model-based	Probabilistic NN
	Dreamer (Hafner et al., 2019a)	First-order	Model-based	Probabilistic NN
	IVG (Byravan et al., 2020)	First-order	Model-free	Deterministic NN
	SAC-SVG (Amos et al., 2021)	First-order	Model-free	Deterministic NN
	BPTT	First-order	-	Differentiable sim.
	SHAC (Xu et al., 2022)	First-order	Model-free	Differentiable sim.
	AHAC (this paper)	First-order	Model-free	Differentiable sim.

Table 2: Comparison between recent influential RL algorithms for continuous control. We classify these approaches into the MFRL, ZO-MBRL and FO-MBRL categories predominantly by the way policy(actor) is learned. Zeroth-order (policy gradient) methods are harnessing the gradient estimates following Equation 2 without the need of taking dynamics model gradients, while First-order methods differentiate the whole trajectory following Equation 3. Model-based Value Learning refers to methods that fall under the Model-based Value Expansion (MVE) approach (Feinberg et al., 2018).

Recent interest in differentiable simulation has given rise to several works that employ FOBG to optimise objectives by back-propagating through the dynamics model (Hu et al., 2019b; Liang et al., 2019; Huang et al., 2021; Du et al., 2021). Although not explicitly addressing RL, these works follow the idea of rolling out a trajectory under a policy and iteratively optimising it until convergence. This approach can be reformulated as a FO-MBRL algorithm, referred to as Back-Propagation-Through-Time (BPTT).

When employed for typical long episodic RL tasks, BPTT performs poorly due to a noisy optimisation landscape and exploding gradients. (Xu et al., 2022) proposes Short Horizon Actor-Critic (SHAC) to address these issues by (1) introducing a model-free critic that acts as a smooth surrogate of the reward-maximisation objective and (2) by employing short rollouts to avoid high and unstable policy gradients. When run in a massively parallel fashion, SHAC stands out as one of the few MBRL approaches that achieves comparable asymptotic performance to MFRL methods while also demonstrating significantly better sample efficiency.

7 Conclusion

Our study aimed to compare the asymptotic performance of conventional Zeroth-order Model-Free RL (MFRL) methods with First-Order Model-Based (FO-MBRL) methods in differentiable simulators. We assessed the difference between both types of gradients and derived Lemma 3.1 showing that first-order batch gradient (FDBG) empirical bias is upper-bounded by the stiffness of the dynamics. Unfortunately, contact-rich tasks exhibit such properties, which translates to FDBG estimates with high bias, leading to unstable learning.

We initially explored this issue in a toy problem and then introduced an algorithm designed to mitigate the accumulation of gradient bias stemming from stiff dynamics by truncating trajectories upon contact. When applied to high-dimensional locomotion tasks, our proposed approach, Adaptive Horizon Actor-Critic (AHAC), achieved up to a 91% increase in asymptotic episodic rewards compared to state-of-the-art MFRL methods while also exhibiting lower variance. Surprisingly, we found that even with near-infinite data, MFRL methods cannot solve tasks with similar performance to our proposed method. Additionally, AHAC retained the advantages commonly observed in FO-MBRL approaches, including exceptional sample efficiency and scalability to higher-dimensional challenges. Notably, our approach demonstrated the ability to learn complex locomotion policies for a quadruped robot in as little as 10 minutes on a single GPU, paving the way for substantial RL scalability.

While AHAC outperforms MFRL methods in asymptotic rewards, it necessitates the development of differentiable simulators, requiring substantial engineering effort. Thus, we cannot help but admire the simple yet capable model-free algorithms such as PPO. Despite this, the performance of our proposed method AHAC renders it promising for robotic applications. However, it’s essential to acknowledge the sim2real gap, which requires further exploration. Our vision for the next phase involves applying FO-MBRL approaches to real robots in a closed-loop manner where simulations aid policy learning but continually adapt to match the real environment. Furthermore, we believe that our proposed approach, AHAC, still has room for improvement, in particular by truncating each parallel environment by contact stiffness precisely instead of using a uniform rollout horizon.

References

- Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, et al. Solving rubik’s cube with a robot hand. *arXiv preprint arXiv:1910.07113*, 2019.
- Brandon Amos, Samuel Stanton, Denis Yarats, and Andrew Gordon Wilson. On the model-based stochastic value gradient for continuous reinforcement learning. In *Learning for Dynamics and Control*, pp. 6–20. PMLR, 2021.
- Albert S Berahas, Liyuan Cao, Krzysztof Choromanski, and Katya Scheinberg. A theoretical and empirical comparison of gradient approximations in derivative-free optimization. *Foundations of Computational Mathematics*, 22(2):507–560, 2022.
- Jacob Buckman, Danijar Hafner, George Tucker, Eugene Brevdo, and Honglak Lee. Sample-efficient reinforcement learning with stochastic ensemble value expansion. *Advances in neural information processing systems*, 31, 2018.
- Arunkumar Byravan, Jost Tobias Springenberg, Abbas Abdolmaleki, Roland Hafner, Michael Neunert, Thomas Lampe, Noah Siegel, Nicolas Heess, and Martin Riedmiller. Imagined value gradients: Model-based policy optimization with transferable latent dynamics models. In *Conference on Robot Learning*, pp. 566–589. PMLR, 2020.
- Tao Du, Kui Wu, Pingchuan Ma, Sebastien Wah, Andrew Spielberg, Daniela Rus, and Wojciech Matusik. Diffpd: Differentiable projective dynamics. *ACM Transactions on Graphics (TOG)*, 41(2):1–21, 2021.
- John C Duchi, Peter L Bartlett, and Martin J Wainwright. Randomized smoothing for stochastic optimization. *SIAM Journal on Optimization*, 22(2):674–701, 2012.
- Vladimir Feinberg, Alvin Wan, Ion Stoica, Michael I Jordan, Joseph E Gonzalez, and Sergey Levine. Model-based value expansion for efficient model-free reinforcement learning. In *Proceedings of the 35th International Conference on Machine Learning (ICML 2018)*, 2018.
- C Daniel Freeman, Erik Frey, Anton Raichuk, Sertan Girgin, Igor Mordatch, and Olivier Bachem. Brax—a differentiable physics engine for large scale rigid body simulation. *arXiv preprint arXiv:2106.13281*, 2021.
- Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.
- Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019a.
- Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning latent dynamics for planning from pixels. In *International conference on machine learning*, pp. 2555–2565. PMLR, 2019b.
- Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse domains through world models. *arXiv preprint arXiv:2301.04104*, 2023.
- Eric Heiden, Miles Macklin, Yashraj S Narang, Dieter Fox, Animesh Garg, and Fabio Ramos. DiSECt: A Differentiable Simulation Engine for Autonomous Robotic Cutting. *Robotics: Science and Systems*, 2021.
- Yuanming Hu, Luke Anderson, Tzu-Mao Li, Qi Sun, Nathan Carr, Jonathan Ragan-Kelley, and Frédo Durand. DiffTaichi: Differentiable programming for physical simulation. *arXiv preprint arXiv:1910.00935*, 2019a.
- Yuanming Hu, Jiancheng Liu, Andrew Spielberg, Joshua B Tenenbaum, William T Freeman, Jiajun Wu, Daniela Rus, and Wojciech Matusik. Chainqueen: A real-time differentiable physical simulator for soft robotics. In *2019 International conference on robotics and automation (ICRA)*, pp. 6265–6271. IEEE, 2019b.

- Zhiao Huang, Yuanming Hu, Tao Du, Siyuan Zhou, Hao Su, Joshua B Tenenbaum, and Chuang Gan. Plasticinelab: A soft-body manipulation benchmark with differentiable physics. *arXiv preprint arXiv:2104.03311*, 2021.
- Marco Hutter, Christian Gehring, Dominic Jud, Andreas Lauber, C Dario Bellicoso, Vassilios Tsounis, Jemin Hwangbo, Karen Bodie, Peter Fankhauser, Michael Bloesch, et al. Anymal-a highly mobile and dynamic quadrupedal robot. In *2016 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pp. 38–44. IEEE, 2016.
- Jemin Hwangbo, Inkyu Sa, Roland Siegwart, and Marco Hutter. Control of a quadrotor with reinforcement learning. *IEEE Robotics and Automation Letters*, 2(4):2096–2103, 2017.
- Jemin Hwangbo, Joonho Lee, Alexey Dosovitskiy, Dario Bellicoso, Vassilios Tsounis, Vladlen Koltun, and Marco Hutter. Learning agile and dynamic motor skills for legged robots. *Science Robotics*, 4(26):eaau5872, 2019.
- Michael Janner, Justin Fu, Marvin Zhang, and Sergey Levine. When to trust your model: Model-based policy optimization. In *Advances in Neural Information Processing Systems*, 2019.
- Juraj Kabzan, Lukas Hewing, Alexander Liniger, and Melanie N Zeilinger. Learning-based model predictive control for autonomous racing. *IEEE Robotics and Automation Letters*, 4(4):3363–3370, 2019.
- Elia Kaufmann, Antonio Loquercio, René Ranftl, Matthias Müller, Vladlen Koltun, and Davide Scaramuzza. Deep drone acrobatics. *arXiv preprint arXiv:2006.05768*, 2020.
- Vijay Konda and John Tsitsiklis. Actor-critic algorithms. *Advances in neural information processing systems*, 12, 1999.
- Ashish Kumar, Zipeng Fu, Deepak Pathak, and Jitendra Malik. Rma: Rapid motor adaptation for legged robots. *arXiv preprint arXiv:2107.04034*, 2021.
- Junbang Liang, Ming Lin, and Vladlen Koltun. Differentiable cloth simulation for inverse problems. *Advances in Neural Information Processing Systems*, 32, 2019.
- Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Nikita Rudin, David Hoeller, Philipp Reist, and Marco Hutter. Learning to walk in minutes using massively parallel deep reinforcement learning. In *Conference on Robot Learning*, pp. 91–100. PMLR, 2022.
- John Schulman, Nicolas Heess, Theophane Weber, and Pieter Abbeel. Gradient estimation using stochastic computation graphs. *Advances in neural information processing systems*, 28, 2015.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- Hyung Ju Suh, Max Simchowitz, Kaiqing Zhang, and Russ Tedrake. Do differentiable simulators give better policy gradients? In *International Conference on Machine Learning*, pp. 20668–20696. PMLR, 2022.
- Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. *Advances in neural information processing systems*, 12, 1999.

- Joel A Tropp et al. An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning*, 8(1-2):1–230, 2015.
- Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Reinforcement learning*, pp. 5–32, 1992.
- Jie Xu, Tao Chen, Lara Zlokapa, Michael Foshey, Wojciech Matusik, Shinjiro Sueda, and Pulkit Agrawal. An end-to-end differentiable framework for contact-aware robot design. *arXiv preprint arXiv:2107.07501*, 2021.
- Jie Xu, Viktor Makoviychuk, Yashraj Narang, Fabio Ramos, Wojciech Matusik, Animesh Garg, and Miles Macklin. Accelerated policy learning with parallel differentiable simulation. *arXiv preprint arXiv:2204.07137*, 2022.

A Proof of Lemma 3.1

Assumption A.1. As well as being continuously differentiable (Assumption 2.7, the policy is 1-Lipshitz-smooth: $\|\nabla_{\theta}\pi_{\theta}(\mathbf{a}|\mathbf{s})\| \leq B_{\pi} \leq 1$ and the reward function is 1 Lipshitz-smooth and bounded rewards $r(\mathbf{s}, \mathbf{a}) \leq \|\nabla r(\mathbf{s}, \mathbf{a})\| \leq B_r \leq 1 \ \forall \mathbf{s} \in \mathbb{R}^n; \mathbf{a} \in \mathbb{R}^m; \theta \in \mathbb{R}^d$.

Proof. First, we expand our definition of bias and define a random variable of a single Monte-Carlo sample

$$\begin{aligned} \left\| \text{Var}[\nabla]_{\theta}^{[1]} J(\theta) - \text{Var}[\nabla]_{\theta}^{[0]} J(\theta) \right\| &= \left\| \frac{1}{N} \sum_{i=1}^N \hat{\nabla}_{\theta}^{[1]} J_i(\theta) - \frac{1}{N} \sum_{i=1}^N \hat{\nabla}_{\theta}^{[0]} J_i(\theta) \right\| \\ &= \frac{1}{N} \left\| \sum_{i=1}^N (\hat{\nabla}_{\theta}^{[1]} J_i(\theta) - \hat{\nabla}_{\theta}^{[0]} J_i(\theta)) \right\| \end{aligned} \quad (10)$$

Define $X_i = \hat{\nabla}_{\theta}^{[1]} J_i(\theta) - \hat{\nabla}_{\theta}^{[0]} J_i(\theta)$ and bound it:

$$\begin{aligned} X_i &= \sum_{h=1}^H \left(\nabla_{\mathbf{a}_h} r(\mathbf{s}_h, \mathbf{a}_h) \nabla_{\theta} \pi_{\theta}(\mathbf{a}_h | \mathbf{s}_h) + \sum_{h'=1}^{h-1} \nabla_{\mathbf{s}_h} r(\mathbf{s}_h, \mathbf{a}_h) \left(\prod_{t=1}^{h'} \nabla_{\mathbf{s}_t} f(\mathbf{s}_t, \mathbf{a}_t) \right) \nabla_{\theta} \pi_{\theta}(\mathbf{a}_{h'} | \mathbf{s}_{h'}) \right) \\ &\quad + \sum_{h=1}^H r(\mathbf{s}_h, \mathbf{a}_h) \nabla \log \pi_{\theta}(\mathbf{a}_h | \mathbf{s}_h) \\ &= \sum_{h=1}^H \left(\nabla_{\mathbf{a}_h} r(\mathbf{s}_h, \mathbf{a}_h) \nabla_{\theta} \pi_{\theta}(\mathbf{a}_h | \mathbf{s}_h) - r(\mathbf{s}_h, \mathbf{a}_h) \nabla_{\theta} \log \pi_{\theta}(\mathbf{a}_h | \mathbf{s}_h) \right. \\ &\quad \left. + \sum_{h'=1}^{h-1} \nabla_{\mathbf{s}_h} r(\mathbf{s}_h, \mathbf{a}_h) \left(\prod_{t=1}^{h'} \nabla_{\mathbf{s}_t} f(\mathbf{s}_t, \mathbf{a}_t) \right) \nabla_{\theta} \pi_{\theta}(\mathbf{a}_{h'} | \mathbf{s}_{h'}) \right) \\ &\leq \sum_{h=1}^H \sum_{h'=1}^{h-1} \nabla_{\mathbf{s}_h} r(\mathbf{s}_h, \mathbf{a}_h) \left(\prod_{t=1}^{h'} \nabla_{\mathbf{s}_t} f(\mathbf{s}_t, \mathbf{a}_t) \right) \nabla_{\theta} \pi(\cdot, \mathbf{s}_{h'}) \quad (\text{Assumption A.1}) \\ &\leq \sum_{h=1}^H \sum_{h'=1}^{h-1} B_r B_{\pi} \prod_{t=1}^{h'} \|\nabla f(\mathbf{s}_t, \mathbf{a}_t)\| \end{aligned}$$

We can now sum up these random variables as $Z = \sum_{i=1}^N X_i$ and create an upper concentration bound. As these RVs are difficult to bound, we can apply a Chebyshev Inequality (Tropp et al., 2015):

$$P(\|Z - \mathbb{E}[Z]\| > \epsilon) \leq \frac{\text{Var}[Z]}{\epsilon^2}$$

Since the gradient samples are assumed to be i.i.d., we can expand this variance using the definition of each gradient type where all expectations are taken over the action distributions \mathbf{a}_h for each step:

$$\begin{aligned}
\text{Var}[Z] &= \text{Var}\left[\sum_{i=1}^N X_i\right] = \text{Var}\left[\sum_{i=1}^N \hat{\nabla}_{\boldsymbol{\theta}}^{[1]} J_i(\boldsymbol{\theta}) - \hat{\nabla}_{\boldsymbol{\theta}}^{[0]} J_i(\boldsymbol{\theta})\right] \\
&= \sum_{i=1}^N \text{Var}\left[\hat{\nabla}_{\boldsymbol{\theta}}^{[1]} J_i(\boldsymbol{\theta}) - \hat{\nabla}_{\boldsymbol{\theta}}^{[0]} J_i(\boldsymbol{\theta})\right] \\
&\leq \sum_{i=1}^N \mathbb{E}\left[\left\|\hat{\nabla}_{\boldsymbol{\theta}}^{[1]} J_i(\boldsymbol{\theta}) - \hat{\nabla}_{\boldsymbol{\theta}}^{[0]} J_i(\boldsymbol{\theta})\right\|^2\right] \\
&\leq \sum_{i=1}^N \mathbb{E}\left[\left\|\sum_{h=1}^H \sum_{h'=1}^{h-1} B_r B_\pi \prod_{t=1}^{h'} \|\nabla f(\mathbf{s}_t, \mathbf{a}_t)\| \right\|^2\right] \\
&\leq NH^4 B_r^2 B_\pi^2 \mathbb{E}\left[\prod_{t=1}^H \|\nabla f(\mathbf{s}_t, \mathbf{a}_t)\|^2\right]
\end{aligned}$$

With this result we can return back to Equation 10 and obtain

$$\begin{aligned}
\left\|\text{Var}[\nabla_{\boldsymbol{\theta}}^{[1]} J(\boldsymbol{\theta})] - \text{Var}[\nabla_{\boldsymbol{\theta}}^{[0]} J(\boldsymbol{\theta})]\right\| &\leq \frac{1}{N} NH^4 B_r^2 B_\pi^2 \mathbb{E}\left[\prod_{t=1}^H \|\nabla f(\mathbf{s}_t, \mathbf{a}_t)\|^2\right] \\
&= H^4 B_r^2 B_\pi^2 \mathbb{E}\left[\prod_{t=1}^H \|\nabla f(\mathbf{s}_t, \mathbf{a}_t)\|^2\right]
\end{aligned}$$

□

B Summary of modifications

To develop Adaptive Horizon Actor Critic (AHAC) algorithm, we used the Short Horizon Actor Critic (SHAC) algorithm (Xu et al., 2022) as a starting point. This section details all modifications applied to the SHAC in order to derive AHAC and achieve the reported results in this paper. We also note that some of these are not exclusive to either approach.

1. **Adaptive horizon objective** - instead of optimising for the short horizon rollout return, we introduce the new constrained objective shown in Equation 8. To optimise that and adapt the horizon H , we introduced the dual problem in Equation 9 and optimised it directly for policy parameters $\boldsymbol{\theta}$ and the Lagrangian coefficients ϕ .

$$\begin{aligned}
J(\boldsymbol{\theta}) &:= \underbrace{\sum_{h=t}^{t+T-1} \gamma^{h-t} r(\mathbf{s}_h, \mathbf{a}_h) + \gamma^t V_\psi(\mathbf{s}_{t+T})}_{\text{SHAC objective}} \quad \text{s.t.} \quad \underbrace{\|\nabla f(\mathbf{s}_t, \mathbf{a}_t)\| \leq C \quad \forall t \in \{0, \dots, T\}}_{\text{AHAC objective}}
\end{aligned}$$

2. **Clipped dynamics gradient norms** - SHAC clips the full gradient norm of Equation 7 i.e. $gc(\nabla_{\boldsymbol{\theta}} J(\boldsymbol{\theta}))$ similar to commonly found PPO implementations (Schulman et al., 2017). On the other hand, AHAC takes inspiration from our proposed Lemma 3.1 and applies the gradient norm clipping on every step of the H -timestep rollout:

$$\begin{aligned}
\tilde{\nabla}_{\mathbf{s}_h} f(\mathbf{s}_h, \mathbf{a}_h) &= gc(\nabla_{\mathbf{s}_h} f(\mathbf{s}_h, \mathbf{a}_h)) \quad \forall h \in [0, H] \\
\tilde{\nabla}_{\mathbf{a}_h} f(\mathbf{s}_h, \mathbf{a}_h) &= gc(\nabla_{\mathbf{a}_h} f(\mathbf{s}_h, \mathbf{a}_h)) \quad \forall h \in [0, H]
\end{aligned}
\quad gc(\mathbf{x}) := \begin{cases} \frac{\mathbf{x}}{\|\mathbf{x}\|_2} & \text{if } \|\mathbf{x}\|_2 > 1 \\ \mathbf{x} & \text{otherwise} \end{cases}$$

3. **Dual critic** - the original implementation of SHAC struggled with more complex tasks such as Humanoid due to its highly non-convex value landscape. The authors of (Xu et al., 2022) solved that by introducing a delayed target critic similar to prior work in deep RL (Lillicrap et al., 2015). We found that approach brittle and requiring more hyper-parameter tuning. Instead, we replaced it with a dual critic similar to (Haarnoja et al., 2018) which has been shown to stabilise on-policy algorithms (Amos et al., 2021). For our work, we found that it reduced variance of asymptotic rewards achieved by AHAC while removing a hyperparameter.
4. **Critic training until convergence** - empirically we found that different problems present different value landscapes. The more complex the landscape, the more training the critic required and the critic often failed to fit the data with the limited number of critic iterations done in SHAC (16). Instead of training the critic for a fixed number of iterations, we trained the (dual) critic of AHAC until convergence defined by $\sum_{i=n-5}^n \mathcal{L}_i(\psi) - \mathcal{L}_{i-1}(\psi) < 0.5$ where $\mathcal{L}_i(\psi)$ is the critic loss for mini-batch iteration i . We allowed the critic to be trained for a maximum of 64 iterations. We found that this resulted in asymptotic performance improvements on more complex tasks such as Humanoid and SNU Humanoid while removing yet another hyper-parameter.

C AHAC-1 algorithm

Algorithm 2: Adaptive Horizon Actor-Critic (Single environment)

```

1 Given:  $\gamma$ : discount rate
2 Given:  $\alpha$ : learning rate
3 Given:  $H$ : maximum trajectory length
4 Given:  $C$ : contact threshold
5 Initialise learnable parameters  $\theta, \psi$ 
6  $t \leftarrow 0$ 
7 while episode not done do
    /* rollout policy */
8     Initialise buffer  $D$ 
9     Initialise return  $R \leftarrow 0$ 
10    while  $\|\nabla f\| \leq C$  and  $h \leq H$  do
11         $\mathbf{a}_t \sim \pi_{\theta}(\cdot | \mathbf{s}_t)$ 
12         $\mathbf{s}_{t+1} = f(\mathbf{s}_t, \mathbf{a}_t)$ 
13         $D \leftarrow D \cup \{(\mathbf{s}_{t+h}, \mathbf{a}_{t+h}, \mathbf{r}_{t+h}, V_{\psi}(\mathbf{s}_{t+h+1}))\}$ 
14         $t \leftarrow t + 1$ 
    /* train actor with Eq. 7 */
15     $\theta \leftarrow \theta - \alpha \nabla_{\theta} J(\theta)$ 
    /* train critic with Eq. 6 */
16    while not converged do
17        sample  $(\mathbf{s}, \hat{V}(\mathbf{s})) \sim D$ 
18         $\psi \leftarrow \psi + \alpha \nabla_{\psi} \mathcal{L}(\psi)$ 

```

D Hyper-parameters

This section details all hyper-parameters used in the main experiments of Section 5. Table 3 shows common hyper-parameters shared between all tasks. While table 4 shows hyper-parameters specific to each problem. Where possible we attempted to use the hyper-parameters suggested by the original works, however, we also attempted to share hyper-parameters between algorithms to ease comparison. If a specific hyper-parameter is not mentioned, then it is the one used in the original work behind the specific algorithm.

	AHAC	SHAC	PPO	SAC	SVG
Mini-epochs		16	5		4
Batch size	8	8	8	32	1024
λ	0.95	0.95	0.95		
γ	0.99	0.99	0.99	0.99	0.99
H - horizon		32	32		3
C - contact thresh.	500				
Grad norm	1.0	1.0	1.0		
ϵ			0.2		
Actor $\log(\sigma)$ bounds				(-5,2)	(-5,2)
α - temperature				0.2	0.1
λ_α				10^{-4}	10^{-4}
$ D $ - buffer size				10^6	10^6
Seed steps	0	0	0	10^4	10^4

Table 3: Table of hyper-parameters for all algorithms benchmarked in Section 5. These are shared across all tasks.

	Hopper	Ant	Anymal	Humanoid	SNU Humanoid
Actor layers	(128, 64, 32)	(128, 64, 32)	(256, 128)	(256, 128)	(512, 256)
Actor α_θ	2×10^{-3}	2×10^{-3}	2×10^{-3}	2×10^{-3}	2×10^{-3}
Critic layers	(64, 64)	(64, 64)	(256, 128)	(256, 128)	(256, 256)
Critic α_ψ	4×10^{-3}	2×10^{-3}	2×10^{-3}	5×10^{-4}	5×10^{-4}
Critic τ	0.2	0.2	0.2	0.995	0.995

Table 4: Task-specific hyper-parameters. All benchmarked algorithms share the same actor and critic network hyper-parameters with ELU activation functions. AHAC and PPO do not have target critic networks and as such do not have τ as a hyper-parameter.