
How to Learn a Useful Critic? Model-based Action-Gradient-Estimator Policy Optimization

Pierluca D’Oro*
MILA, Université de Montréal
pierluca.doro@mila.quebec

Wojciech Jaśkowski
NNAISENSE
wojciech@nnaisense.com

Abstract

Deterministic-policy actor-critic algorithms for continuous control improve the actor by plugging its actions into the critic and ascending the action-value gradient, which is obtained by chaining the actor’s Jacobian matrix with the gradient of the critic with respect to input actions. However, instead of gradients, the critic is, typically, only trained to accurately predict expected returns, which, on their own, are useless for policy optimization. In this paper, we propose MAGE, a model-based actor-critic algorithm, grounded in the theory of policy gradients, which explicitly learns the action-value gradient. MAGE backpropagates through the learned dynamics to compute gradient targets in temporal difference learning, leading to a critic tailored for policy improvement. On a set of MuJoCo continuous-control tasks, we demonstrate the efficiency of the algorithm in comparison to model-free and model-based state-of-the-art baselines.

1 Introduction

Reinforcement learning (RL) [36, 52] studies sequential decision making problems, in which an agent aims at maximizing the cumulative reward it collects in an environment. One of the most popular classes of algorithms for RL are *policy gradient methods* [10, 53], which involve differentiable control policies improved by gradient ascent. They feature suitability to environments with continuous state and action spaces, and compatibility with state-of-the-art deep learning [42] methods. Policy gradient algorithms often employ an actor-critic [26] scheme: an *actor*, which determines the control policy, is evaluated using a *critic*. Thus, the degree of actor’s improvement is limited by the information provided by the critic, naturally raising the question of how the critic should be trained.

Typically, algorithms that use powerful function approximators [18, 28] learn the critic by temporal difference [50], optimizing for an accurate prediction of the expected return of the actor. For deterministic-policy continuous-control [28, 47], however, the value provided by the critic is neither used for improving the policy nor for acting in the environment [53]. Instead, only the *action-gradient* of the value function, i.e., the gradient of the critic w.r.t. the action performed by the actor, is employed during policy optimization. Specifically, the policy gradient is obtained through the computation of the *action-value gradient*, by chaining the actor’s Jacobian with the action-gradient of the critic.

Learning the critic *by value* rather than *by action-gradient* of the value relies on hazy smoothness assumptions on the real value function [47]. This means that, in conventional temporal difference learning, the critic learns action-value gradients *implicitly*, which could harm the performance of a deterministic policy gradient algorithm.

In this paper, we propose Model-based Action-Gradient-Estimator Policy Optimization (MAGE), a continuous-control deterministic-policy actor-critic algorithm that *explicitly* trains the critic to provide accurate action-gradients for the use in the policy improvement step. Motivated by both the theory on

*Work done while at NNAISENSE.

Deterministic Policy Gradients [47] and practical considerations, MAGE uses temporal difference methods to minimize the error on the action-value gradient. For this, the algorithm leverages a trained dynamics model as a proxy for a differentiable environment and techniques reminiscent of double backpropagation [12]. On a challenging continuous control benchmark [6, 55], we show that MAGE is significantly more sample-efficient than state-of-the-art model-free and model-based baselines.

The rest of the paper is organized as follows. In Section 2, we provide the notation and background on deterministic policy gradients. Our algorithm, together with its theoretical motivation, is introduced in Section 3, followed by empirical results in Section 4. In Section 5, we present some of the related work and its relationship with our approach.

2 Background

2.1 Preliminaries

Consider a discrete-time Markov Decision Process [36] (MDP), defined as $\mathcal{M} = (\mathcal{S}, \mathcal{A}, p, r, \gamma, \mu)$, where \mathcal{S} is the space of possible states, \mathcal{A} is the space of possible actions, $p : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ is the transition model, $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the known and differentiable reward function, γ is the discount factor, $\mu \in \Delta(\mathcal{S})$ is the initial state distribution. The behavior of the agent is described by a deterministic policy $\pi_\theta : \mathcal{S} \rightarrow \mathcal{A}$, belonging to a parametric space of policies $\Pi = \{\pi_\theta : \theta \in \Theta \subseteq \mathbb{R}^n\}$, for which we will occasionally omit the parameter subscript. Let d_μ^π be the γ -discounted state distribution induced by policy π_θ , defined as $d_\mu^\pi(s) = (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \Pr(s_t = s | \pi, \mu)$. The total reward collected by an agent is quantified with action-value function $Q^\pi(s, a) = \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) | s_0 = s, a_0 = a]$ and performance function $J(\theta) = \mathbb{E}_{s \sim \mu} [Q^\pi(s, \pi_\theta(s))]$.

Practical algorithms can employ an approximate action-value function \hat{Q} and an approximate dynamics model \hat{p} , which, most commonly, are parametric function approximators specified by the spaces $\mathcal{Q} = \{Q_\phi : \phi \in \Phi \subseteq \mathbb{R}^h\}$ and $\mathcal{P} = \{p_\omega : \omega \in \Omega \subseteq \mathbb{R}^k\}$.

2.2 Deterministic Policy Gradients and TD-learning

Policy gradient methods improve the policy π_θ by ascending the direction of the gradient of its performance function $J(\theta)$. The *Deterministic Policy Gradient Theorem* [47] provides a practical way to calculate this gradient. It shows that, under some mild regularity conditions on the MDP, the gradient of the performance of a deterministic policy π_θ is given by:

$$\nabla_\theta J(\theta) = \frac{1}{1 - \gamma} \int_{\mathcal{S}} d_\mu^\pi(s) \nabla_a Q^\pi(s, a) \Big|_{a=\pi_\theta(s)} \nabla_\theta \pi_\theta(s) ds. \quad (1)$$

This result can be interpreted through the lens of the chain rule applied to the *action-value gradient* $\nabla_\theta Q^\pi$: the policy gradient does not directly depend on the gradient of d_μ^π , and can be obtained by just chaining the actor’s Jacobian $\nabla_\theta \pi_\theta$ with the *action-gradient* of the value function $\nabla_a Q^\pi$.

The theorem motivates a family of policy gradient actor-critic algorithms, such as DDPG [28] and TD3 [18]. Similarly to the classical policy iteration [52], the evaluation of a policy $\pi \in \Pi$ (called *actor* in this context) is interleaved with its improvement w.r.t the approximate action-value function $\hat{Q} \in \mathcal{Q}$ (called *critic*). Specifically, the typical desideratum consists in finding a critic \hat{Q} which minimizes the *policy evaluation error*:

$$\hat{Q} \in \arg \min_{\tilde{Q} \in \mathcal{Q}} \mathbb{E}_{s \sim d_\mu^\pi} \left| \delta^{\pi, \tilde{Q}}(s, \pi(s)) \right|, \quad (2)$$

where $\delta^{\pi, \tilde{Q}}(s, a) = Q^\pi(s, a) - \tilde{Q}(s, a)$ is a deviation w.r.t the true state-action value. Given the lack of knowledge about the transition model, Q^π needs to be approximated. A common approximation technique consists in employing the *temporal difference* (TD) error [50], defined as $\hat{\delta}^{\pi, \tilde{Q}}(s, a, s') = r(s, a) + \gamma \tilde{Q}(s', \pi(s')) - \tilde{Q}(s, a)$, giving rise to a bootstrapped optimization criterion for \hat{Q} :

$$\hat{Q} \in \arg \min_{\tilde{Q} \in \mathcal{Q}} \mathbb{E}_{\substack{s \sim d_\mu^\pi \\ s' \sim p(\cdot | s, \pi(s))}} \left| \hat{\delta}^{\pi, \tilde{Q}}(s, \pi(s), s') \right|. \quad (3)$$

Minimizing the TD-error, albeit under rather strong assumptions, enjoys convergence guarantees [52, 56]. Once a critic is learned, the actor π_θ can be improved by maximizing the action-value function for actions produced by the policy:

$$\pi_\theta \in \arg \max_{\tilde{\pi}_\theta \in \Pi_\Theta} \mathbb{E}_{s \sim d_\mu^\pi} \left[\widehat{Q}(s, \tilde{\pi}_\theta(s)) \right]. \quad (4)$$

The above can be seen as a generalization of the policy improvement step in classical policy iteration, which relies on maximization over a discrete action space that cannot be easily carried out in continuous spaces. In practice, to reduce the computational burden, the problems in Equation 3 and Equation 4 are solved only partially (e.g., by using a single optimization step) at each iteration, similarly to generalized policy iteration [52].

3 Learning Action-Value Gradients

In this section, we discuss theoretically how to learn a useful critic in the context of deterministic policy gradients. Then, we make the theoretical insights concrete and, guided by practical considerations, present *Model-based Action-Gradient-Estimator Policy Optimization* (MAGE), a novel policy optimization algorithm.

3.1 How to Learn a Useful Critic?

An actor can only be as good as allowed by its critic. Thus, obtaining an *effective* critic is one of the most crucial passages for any actor-critic algorithm. In the previous section, we outlined the most common method to train the critic, consisting in the minimization of the temporal difference error. However, when the learned action-value function will not be perfect, as common in policy optimization with function approximation, minimizing the TD-error does not guarantee that the critic will be effective for the goal of solving the control problem. Instead, the following result provides foundations for a more grounded objective function for critic learning.

Proposition 3.1. *Let Π be a parametric space of L_π -Lipschitz continuous differentiable deterministic policies, \mathcal{Q} a space of approximate value functions and $\|\cdot\|$ any p -norm. Given $\pi \in \Pi$ and $\widehat{Q} \in \mathcal{Q}$, the norm of the difference between the true policy gradient $\nabla_\theta J(\theta)$ and its approximation $\widehat{\nabla}_\theta J(\theta)$, which uses \widehat{Q} , can be upper bounded as:*

$$\|\nabla_\theta J(\theta) - \widehat{\nabla}_\theta J(\theta)\| \leq \frac{L_\pi}{1 - \gamma} \mathbb{E}_{s \sim d_\mu^\pi} \left\| \nabla_a \delta^{\pi, \widehat{Q}}(s, a) \Big|_{a=\pi(s)} \right\|.$$

The proposition (see Appendix A for the proof) is a direct consequence of the *Deterministic Policy Gradient Theorem* and is thus valid when deterministic policies are employed. The Lipschitz assumption for π is easily satisfied for many policy classes of practical use, e.g., neural networks [16].

Proposition 3.1 suggests that it is the norm of the action-gradient of the policy evaluation error instead of its value that should be minimized to reduce the bias introduced by the use of the approximate value function \widehat{Q} . To minimize the bound, a proxy for the unknown Q^π is needed. To this aim, it is possible to follow the approach of traditional TD-learning, substituting the evaluation error $\delta^{\pi, \widehat{Q}}$ with the TD-error $\widehat{\delta}^{\pi, \widehat{Q}}$. This leads to the following optimization problem:

$$\widehat{Q} \in \arg \min_{\widehat{Q} \in \mathcal{Q}} \mathbb{E}_{\substack{s \sim d_\mu^\pi \\ s' \sim p(\cdot|s, \pi(s))}} \left\| \nabla_a \widehat{\delta}^{\pi, \widehat{Q}}(s, \pi(s), s') \right\|. \quad (5)$$

Notice that computing the gradient w.r.t. the action of the TD-error $\widehat{\delta}^{\pi, \widehat{Q}}$ requires taking into account the effect of action a on the transition to the subsequent state in the environment s' , i.e., backpropagating through the environment dynamics p . Since p is not available in typical RL settings, especially in a differentiable form, it needs to be substituted with an approximate model \widehat{p} , as commonly done in model-based RL [7, 10, 23]. An environment model gives rise to imaginary transitions $(s, \pi(s), \widehat{s})$, where $\widehat{s} \sim \widehat{p}(\cdot|s, \pi(s))$. Given differentiable model, policy, and action-value function, the action-gradient can be effectively computed by leveraging standard automatic differentiation

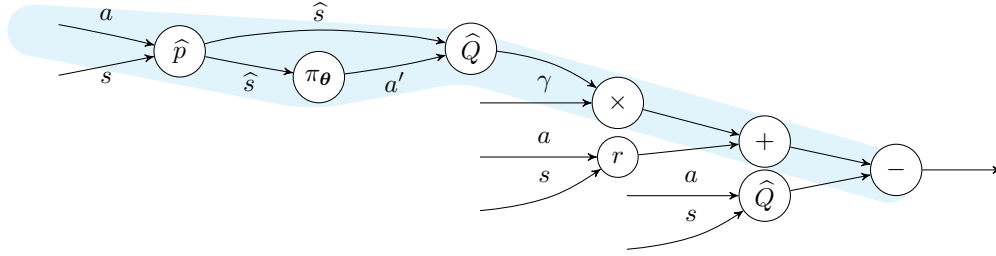


Figure 1: Graph describing the computation of $\widehat{\delta}^{\pi_\theta, \widehat{Q}}$, when using policy π_θ , model \widehat{p} , action-value function \widehat{Q} . Nodes and edges represent functions and variables, respectively. To compute $\nabla_a \widehat{\delta}^{\pi_\theta, \widehat{Q}}$, all the paths from the output back to a must be considered, including the one highlighted in cyan, which involves the environment dynamics. Therefore, an approximate differentiable model \widehat{p} needs to be learned in order to make all the required paths accessible.

tools [5]. The corresponding computational graph is depicted in Figure 1. This leads to a viable way to obtain \widehat{Q} :

$$\widehat{Q} \in \arg \min_{\widehat{Q} \in \mathcal{Q}} \mathbb{E}_{\substack{s \sim d_\mu^\pi \\ \widehat{s} \sim \widehat{p}(\cdot|s, \pi(s))}} \left\| \nabla_a \widehat{\delta}^{\pi, \widehat{Q}}(s, \pi(s), \widehat{s}) \right\|. \quad (6)$$

Even in the general case of a stochastic model, differentiating through the resulting computations is still possible for many commonly used model classes via the reparametrization trick [21]. Using an approximate model \widehat{p} implies a tradeoff, since additional bias is injected into the estimation of the critic. Nonetheless, the use of \widehat{p} is the most direct way to solve the optimization problem in Equation 6 and to obtain a \widehat{Q} that provides a more accurate policy gradient w.r.t. the typical critic.

3.2 Model-based Action-Gradient-Estimator Policy Optimization

The outlined procedure for learning the value function requires an approximate model p_ω , thus naturally suggesting its integration into a model-based policy optimization framework. A model-based actor-critic method involves three steps during each iteration: learning the model p_ω , updating the action-value function Q_ϕ and improving the policy π_θ . In the following, we consider neural networks as function approximators to represent the three modules, although any class of differentiable models could be leveraged. Our approach is inspired by Dyna [51], and employs an approximate dynamics model for generating 1-step imaginary on-policy transitions starting from observed states stored in a replay buffer. Those transitions are then employed to learn Q_ϕ , and, in turn, leveraged for computing an improvement direction for the parameters of the policy π_θ .

In preliminary experiments, we found that directly solving the minimization problem in Equation 6 is hard in practice. During the optimization, the parameters are prone to be trapped in local-minima, which leads to degenerate solutions. A demonstration of this phenomenon is detailed in Appendix B.1. The root cause of this effect is unknown and suggests the existence of a tradeoff between the easier minimization of the TD-error and the more theoretically grounded minimization of its action-gradient.

We propose as a remedy the introduction of a constraint into the optimization problem. We argue that, among the possible solutions, a natural one is constraining the optimization landscape by bounding the traditional TD-error (see Equation 3), and thus solving the following optimization problem:

$$\begin{aligned} \min_{\widehat{\phi} \in \Phi} \quad & \mathbb{E}_{\substack{s \sim d_\mu^\pi \\ \widehat{s} \sim p_\omega(\cdot|s, \pi(s))}} \left\| \nabla_a \widehat{\delta}^{\pi, Q_{\widehat{\phi}}}(s, a, \widehat{s}) \Big|_{a=\pi(s)} \right\| \\ \text{s.t.} \quad & \mathbb{E}_{\substack{s \sim d_\mu^\pi \\ \widehat{s} \sim p_\omega(\cdot|s, \pi(s))}} \left| \widehat{\delta}^{\pi, Q_{\widehat{\phi}}}(s, \pi(s), \widehat{s}) \right| \leq \lambda. \end{aligned} \quad (7)$$

As the above expressions already require non-trivial gradient computations, we avoid the use of complex and expensive methods for nonlinear programming. Instead, we resort to *penalty function*

Algorithm 1 Model-based Action-Gradient-Estimator Policy Optimization (MAGE)

Input: Initial buffer \mathcal{B} , set of parameter vectors $\{\omega, \phi, \theta\}$

for each iteration **do**

 Collect transition (s, a, s') acting according to exploratory version of π_θ

$\mathcal{B} \leftarrow \mathcal{B} \cup \{(s, a, s')\}$

for each model learning step **do**

$\omega \leftarrow \omega - \alpha_p \nabla_\omega \ell(s, a, s'; \omega), \quad (s, a, s') \sim \mathcal{B}$

end for

for each policy optimization step **do**

 Extract state s after sampling $(s, \cdot, \cdot) \sim \mathcal{B}$

$\bar{\phi} \leftarrow \phi$

$\hat{\delta}(s, a, \hat{s}; \phi) \leftarrow r(s, a) + \gamma Q_{\bar{\phi}}(\hat{s}, \pi_\theta(\hat{s})) - Q_\phi(s, a), \quad a = \pi_\theta(s), \hat{s} \sim p_\omega(\cdot|s, a)$

$\phi \leftarrow \phi - \alpha_Q \nabla_\phi \left(\left\| \nabla_a \hat{\delta}(s, a, \hat{s}; \phi) \Big|_{a=\pi_\theta(s)} \right\| + \lambda \left| \hat{\delta}(s, a, \hat{s}; \phi) \right| \right)$

$\theta \leftarrow \theta + \alpha_\pi \nabla_\theta Q_\phi(s, \pi_\theta(s))$

end for

end for

methods [49] by regularizing the original objective by using the TD-error. A similar approach has been used in the past in, e.g., Proximal Policy Optimization (PPO, [45]) to approximately solve different constrained optimization problems.

Eventually, the parameters of Q_ϕ are learned by descending the gradient

$$\nabla_\phi \mathcal{L}(s, a, \hat{s}; \phi, \theta, \omega) = \nabla_\phi \left(\left\| \nabla_a \hat{\delta}^{\pi_\theta, Q_\phi}(s, a, \hat{s}) \Big|_{a=\pi_\theta(s)} \right\| + \lambda \left| \hat{\delta}^{\pi_\theta, Q_\phi}(s, a, \hat{s}) \right| \right) \quad (8)$$

on an imaginary transition (s, a, \hat{s}) . This expression requires computing second-order gradients, which would be computationally expensive if computed w.r.t. to the high-dimensional space of parameters Φ of the Q -function. Here, however, the optimization is affordable since the gradients are computed w.r.t., typically low dimensional, actions. Notice also that the computational overhead of the second term in Equation 8 is minimal, since evaluating the TD-error $\hat{\delta}^{\pi_\theta, Q_\phi}(s, a, \hat{s})$ is anyway, when using automatic differentiation, required to compute its gradient.

We plug our critic training method into a model-based Dyna-like algorithm, giving rise to *Model-based Action-Gradient-Estimator Policy Optimization* (MAGE), which is presented² in Algorithm 1. At each iteration, the dynamics model p_ω is trained to maximize the likelihood of the transitions stored in the experience replay buffer \mathcal{B} , or, equivalently, to minimize an appropriate loss function ℓ :

$$\omega \in \arg \min_{\tilde{\omega} \in \Omega} \mathbb{E}_{(s, a, s') \sim \mathcal{B}} [\ell(s, a, s'; \tilde{\omega})]. \quad (9)$$

Then, for one or more steps, the TD-error for the current policy and action-value function is computed, and used together with its action-gradient to update Q_ϕ , which in turn is leveraged to improve π_θ .

4 Experiments

4.1 Sample-Efficient Continuous Control with MAGE

Algorithm settings The general structure of MAGE is compatible with many actor-critic algorithms with deterministic policies. In this experiment, we employ TD3 [18], a popular, state-of-the-art extension to DDPG [28], as a base policy optimization method. This amounts to the addition of target policy smoothing, delayed policy updates, clipped double Q-learning and target functions. We call this version of our algorithm MAGE-TD3³. After each step of environment interaction, we add the collected transition in the replay buffer \mathcal{B} , train the approximate model p_ω , and update critic and actor 10 times. We employ a single value of $\lambda = 0.2$ for all the environments, since we found MAGE to be reasonably robust to the choice of this hyperparameter (see Appendix B). In order to reduce the impact of model bias, MAGE leverages an ensemble of 8 probabilistic Gaussian-output models, trained by maximum likelihood estimation.

²For simplicity of presentation, an abstract version of MAGE is considered in Algorithm 1. Any actor-critic algorithm with deterministic actor can be then used to instantiate MAGE into a practical incarnation.

³The PyTorch [34] implementation, based on [46], is available at <https://github.com/nnaise/mage>.

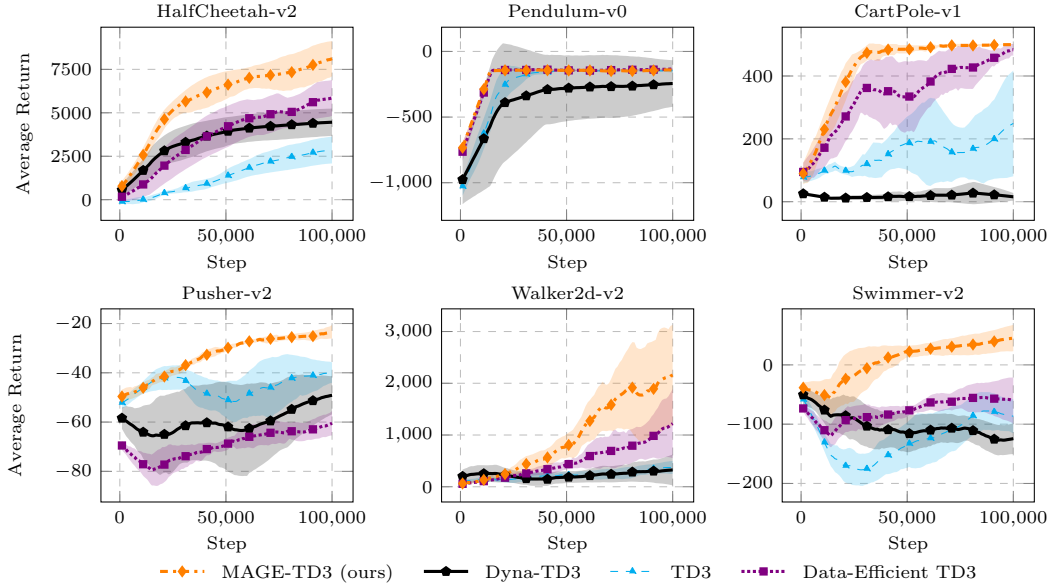


Figure 2: Performance in terms of average return of MAGE on continuous control benchmarks. MAGE compares favorably to the three baselines on all the environments (5 runs, 95% c.i.).

Baselines and environments We consider one model-based and two model-free algorithms as baselines. The first one is Dyna-TD3, which uses a classical TD-error loss, otherwise being identical to MAGE-TD3. It resembles 1-step horizon *Model-based Policy Optimization* (MBPO [23]), but uses a deterministic policy optimized by TD3. Apart from that, we compared MAGE against TD3 and its sample-efficient variant [57], which employs multiple updates for each environment step and trades off computational efficiency and, potentially, stability [32] for sample efficiency. Specifically, for a fair comparison with MAGE-TD3, we execute 10 critic and actor updates after each interaction with the environment. We employ environments from OpenAI Gym [6] and the MuJoCo physics simulator [55] as continuous control benchmarks, assuming, for all the environments, the availability of a differentiable reward function (we will later show that MAGE behaves well also in the absence of this assumption). Additional details concerning the experimental setting are reported in Appendix D.

Results Figure 2 shows the learning curves for the average return of all the approaches. Since our primary interest is MAGE’s sample efficiency, we show the first 10^5 steps of environment interaction. The results show that MAGE is able to learn at least as fast as all the baselines on all the environments, confirming the intuitive advantage of directly optimizing for the accuracy of the estimated action-value gradient. Interestingly, no superiority of the vanilla Dyna-TD3 on its simple data-efficient version can be observed: this demonstrates that there is no intrinsic advantage in terms of sample-efficiency for model-based reinforcement learning, but it is instead highly environment- and algorithm-dependent. On the other hand, increasing the number of offline updates for model-free algorithms can hurt performance in some environments, as it is the case, for instance, on the Pusher-v2 environment. Note that, in contrast with Dyna-TD3, that only leverages the model as a generator for additional transitions w.r.t. the ones that can be obtained in the environment, MAGE makes deeper use of the learned model of the dynamics in order to unlock a peculiar learning modality that would be impossible in a model-free setting. In Appendix B, we also show that MAGE-TD3 matches the asymptotic performance of its model-free counterpart.

4.2 Understanding MAGE

Action-Gradient Estimation MAGE was designed to obtain a critic that is maximally useful for policy improvement by yielding accurate action-value gradients. How much better does it predict them compared to the traditional TD-learning? To investigate this question, we employ the Pendulum-v0 environment, using a differentiable oracle in place of the approximate dynamics model.

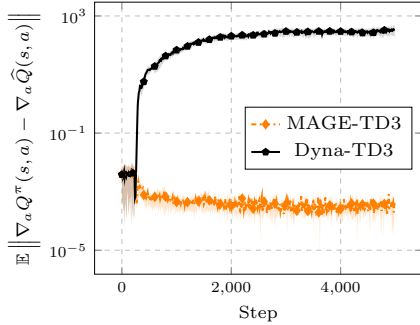


Figure 3: Error of critics in predicting $\nabla_a Q^\pi$ for a random π (4, runs, 95% c.i.). Notice the log scale on the Y axis.

that traditional approaches are able to learn a reasonably good policy even when the learned gradient is very different from the real one is in line with recent analyses on the mismatch between the empirical behavior of policy gradient approaches and their conceptual features [22].

Reward Availability Throughout the presentation and evaluation of MAGE, we assumed complete knowledge of the reward function r of the underlying Markov Decision Process. While this assumption is natural in many real-world settings [10] and thus commonly employed in other model-based reinforcement learning methods [7, 11, 21], its role is particularly crucial in our algorithm. In traditional temporal difference learning, given a transition (s, a, \hat{s}) , the reward $r(s, a)$ constitutes the only grounding element in the objective function. The reward function plays an even stronger role as a grounding element for bootstrapping in MAGE, since both its value $r(s, a)$ and its action-gradient $\nabla_a r(s, a)$ are needed: while the former can be usually observed in the environment, the latter can only be computed with complete knowledge of the underlying function. In our experiments on the sample efficiency of MAGE, we employed the ground-truth reward function (with ground-truth gradients): a natural question is whether MAGE still performs reasonably well if an estimated reward function \hat{r} , learned from data, is used in place of the real r . To answer this question, we evaluate a version of MAGE in which an approximate reward function \hat{r} is learned by using a neural network approximator and minimizing the mean squared error on the rewards observed in the environment. The results, perhaps surprising, are reported in Figure 4 for the Pusher-v2 environment (see Appendix B.2 for the complete results). They show that, for the commonly employed continuous control benchmarks, the performance of our method is only minimally degraded by the use of an approximate reward function in place of the real one, thus suggesting inherent robustness to inaccurate evaluations of the reward function as well as its action-gradient.

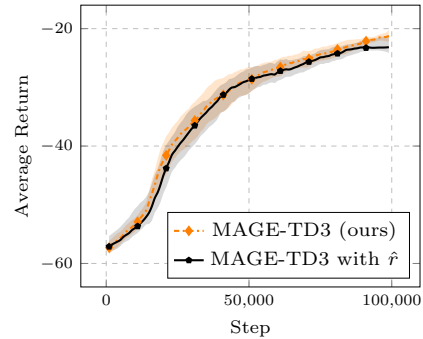


Figure 4: Performance of MAGE on Pusher-v2 using an estimated reward function \hat{r} (5 runs, 95% c.i.).

5 Related Work

Policy gradients are among the most popular methods in reinforcement learning. A variety of algorithms have been proposed for the estimation of the policy gradient, either involving only the policy [4, 32, 61] or also a value function [33, 44, 45]. The latter category of algorithms is referred to as *actor-critic* methods [26, 35]. Among them, the ones based on the *Deterministic Policy Gradient* [28, 47] leverage the action-gradient of the critic. When using function approximation, the quality of the learned critic is of paramount importance [3]: for instance, enforcing on the critic the *compatibility conditions* [47] ensures an unbiased estimate of the policy gradient.

Developed around such conditions, GProp [2] is, to the best of our knowledge, the only method that explicitly optimizes for the accuracy of the learned action-value gradient. It is significantly

different w.r.t. MAGE, being model-free and based on gradient estimation via noisy perturbations together with an additional deviator network. Importantly, while GProp’s deviator network is a function approximator that outputs an estimate for the action-gradient, recent theoretical [39] and practical [40] insights outside of RL suggest that learning the action-gradient by second-order differentiation, as we propose in MAGE, is not only simpler to implement w.r.t. GProp and similar procedures [59], but also fundamentally more effective when using neural network approximators.

The technique we use for learning the action-gradient relies on the differentiation of the TD-error and, thus, of the Bellman equation. This is related to a broad class of methods called *value gradients* [8, 14, 21, 41], in which the policy is improved by backpropagating through the unrolled Bellman equation. Those approaches, however, learn the value function by standard temporal difference [21]. Another classical method, named *Dual Heuristic Programming* (DHP) [13, 35, 60], learns the gradient of the state-value function in a model-based setting, leading to a TD-learning procedure that resembles our approach. However, DHP has the main goal of improving generalization of the value function and exploration, and is fundamentally different from MAGE, that aims at learning an accurate action-gradient of the critic and is motivated by the Deterministic Policy Gradient Theorem.

More broadly, inside and outside of reinforcement learning, several algorithms incorporate gradient penalties into the loss function used for training a neural network. This technique, known as *double backpropagation* [12], has been employed in a number of applications, for instance increasing generalization capabilities [12, 38], enforcing Lipschitz constants [19, 20, 30], or encouraging robustness to adversarial examples [48]. Particularly related to our approach is *Sobolev training* [9], which leverages the availability of the derivatives of a target function to explicitly try to learn both value and gradient of it during supervised training; in our case, no ground-truth gradient is available and we use the action-gradient of the TD-target as a proxy.

Our method learns the action-gradient in the context of model-based policy optimization [10, 54, 58]. We build upon the classical Dyna framework [25, 51], in which a learned model is used for generating imaginary transitions, then employed for training a value function. Our algorithm, which learns a Q-function from model-generated data but only optimizes the policy by using real data, is related to the approaches that compute the policy gradient by using a model-based value function together with trajectories sampled in the environment [1, 11, 21, 23]. In practice, we leverage an ensemble of models, which has been shown to improve performance in a variety of contexts [7, 23, 27].

Finally, our work is related in spirit to *decision-aware model learning* (DAML) [11, 15, 24]. In DAML approaches, the model of the dynamics of the environment is learned by explicitly considering how it will be used for improving the control policy: this is the same rationale behind the learning objective used in MAGE for the critic, focused on how it will be useful for policy optimization, and not merely on how it will be similar to the true value function.

6 Conclusion

In this paper, we presented MAGE, a model-based actor-critic algorithm with deterministic actor, which leverages an approximate dynamics model to directly learn the action-value gradient via temporal difference learning. MAGE employs second-order differentiation to obtain a critic tailored for policy improvement. The empirical evaluation of MAGE demonstrated its superiority over model-based and model-free baselines on challenging high-dimensional continuous control tasks.

A limitation of our method is of computational nature: in addition to the cost of model learning paid also by other model-based actor-critic algorithms, we incur the expense of computing a second-order gradient in order to train the critic, in result, approximately doubling the training time in comparison to the Dyna-based policy gradient approach. This can potentially be alleviated by the development of more efficient automatic differentiation tools, which is, currently, an active area of research [5].

While it is often hard to determine the circumstances under which the addition of an approximate model to a model-free algorithm is beneficial [23], we have shown that model-based techniques, such as MAGE’s gradient-learning procedure, can unlock novel learning modalities, otherwise inaccessible. This can actually be the true power of model-based reinforcement learning. Therefore, apart from improving MAGE (e.g., by investigating the unconstrained critic learning problem) and generalizing it (e.g., to value gradients with real trajectories [21] or multi-steps methods [17]), we hope that future work will reveal other innovative learning schemes that are infeasible in model-free settings.

Broader Impact

The method presented in this paper is a reinforcement learning algorithm that can be used to control a system executing real-valued actions in an environment. Therefore, a natural application of it is in robotics, with positive (e.g., elderly care, resource-efficiency in manufacturing) and negative (e.g., military) applications. Alongside other deep reinforcement learning algorithms, our method is computationally intensive and its training can thus require considerable resources (i.e., hardware and electricity); on the other hand, given that in many real-world scenarios every interaction with a system implies an economic or environmental cost, the sample efficiency of MAGE is aligned with the modern principles of responsible artificial intelligence.

Acknowledgments and Disclosure of Funding

The authors are grateful to Miroslav Štrupl for discussions about the relationship between action-gradients and TD-learning, David Alvarez for co-authoring the reinforcement learning framework used for the experiments, Christian Osendorfer, Miroslav Štrupl, Jan Koutník for their valuable feedback on an early draft of this manuscript, and to everyone at NNAISENSE for contributing to an inspiring research environment.

References

- [1] Pieter Abbeel, Morgan Quigley, and Andrew Y Ng. Using inaccurate models in reinforcement learning. In *Proceedings of the 23rd international conference on Machine learning*, pages 1–8. ACM, 2006.
- [2] David Balduzzi and Muhammad Ghifary. Compatible value gradients for reinforcement learning of continuous deep policies. *arXiv preprint arXiv:1509.03005*, 2015.
- [3] Gabriel Barth-Maron, Matthew W. Hoffman, David Budden, Will Dabney, Dan Horgan, Dhruva TB, Alistair Muldal, Nicolas Heess, and Timothy Lillicrap. Distributional policy gradients. In *International Conference on Learning Representations*, 2018.
- [4] Jonathan Baxter and Peter L Bartlett. Infinite-horizon policy-gradient estimation. *Journal of Artificial Intelligence Research*, 15:319–350, 2001.
- [5] Atilim Gunes Baydin, Barak A. Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind. Automatic differentiation in machine learning: a survey. *Journal of Machine Learning Research*, 18(153):1–43, 2018.
- [6] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym, 2016.
- [7] Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems*, pages 4754–4765, 2018.
- [8] Ignasi Clavera, Yao Fu, and Pieter Abbeel. Model-augmented actor-critic: Backpropagating through paths. In *International Conference on Learning Representations*, 2020.
- [9] Wojciech M Czarnecki, Simon Osindero, Max Jaderberg, Grzegorz Swirszcz, and Razvan Pascanu. Sobolev training for neural networks. In *Advances in Neural Information Processing Systems*, pages 4278–4287, 2017.
- [10] Marc Peter Deisenroth, Gerhard Neumann, Jan Peters, et al. A survey on policy search for robotics. *Foundations and Trends® in Robotics*, 2(1–2):1–142, 2013.
- [11] Pierluca D’Oro, Alberto Maria Metelli, Andrea Tirinzoni, Matteo Papini, and Marcello Restelli. Gradient-aware model-based policy search. In *Thirty-Fourth AAAI Conference on Artificial Intelligence*, 2020.
- [12] Harris Drucker and Yann Le Cun. Improving generalization performance using double back-propagation. *IEEE Transactions on Neural Networks*, 3(6):991–997, 1992.

- [13] Michael Fairbank. Reinforcement learning by value gradients. *ArXiv*, abs/0803.3539, 2008.
- [14] Michael Fairbank. *Value-gradient learning*. PhD thesis, City University London, 2014.
- [15] Amir-massoud Farahmand, Andre Barreto, and Daniel Nikovski. Value-aware loss function for model-based reinforcement learning. In *Artificial Intelligence and Statistics*, pages 1486–1494, 2017.
- [16] Mahyar Fazlyab, Alexander Robey, Hamed Hassani, Manfred Morari, and George Pappas. Efficient and accurate estimation of lipschitz constants for deep neural networks. In *Advances in Neural Information Processing Systems*, pages 11423–11434, 2019.
- [17] Vladimir Feinberg, Alvin Wan, Ion Stoica, Michael I Jordan, Joseph E Gonzalez, and Sergey Levine. Model-based value estimation for efficient model-free reinforcement learning. *arXiv preprint arXiv:1803.00101*, 2018.
- [18] Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In *International Conference on Machine Learning*, pages 1582–1591, 2018.
- [19] Carles Gelada, Saurabh Kumar, Jacob Buckman, Ofir Nachum, and Marc G. Bellemare. Deep-MDP: Learning continuous latent space models for representation learning. In *Proceedings of the 36th International Conference on Machine Learning*, pages 2170–2179, 2019.
- [20] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in neural information processing systems*, pages 5767–5777, 2017.
- [21] Nicolas Heess, Gregory Wayne, David Silver, Timothy Lillicrap, Tom Erez, and Yuval Tassa. Learning continuous control policies by stochastic value gradients. In *Advances in Neural Information Processing Systems*, pages 2944–2952, 2015.
- [22] Andrew Ilyas, Logan Engstrom, Shibani Santurkar, Dimitris Tsipras, Firdaus Janoos, Larry Rudolph, and Aleksander Madry. A closer look at deep policy gradients. In *International Conference on Learning Representations*, 2020.
- [23] 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*, pages 12498–12509, 2019.
- [24] Joshua Joseph, Alborz Geramifard, John W Roberts, Jonathan P How, and Nicholas Roy. Reinforcement learning with misspecified model classes. In *2013 IEEE International Conference on Robotics and Automation*, pages 939–946. IEEE, 2013.
- [25] Gabriel Kalweit and Joschka Boedecker. Uncertainty-driven imagination for continuous deep reinforcement learning. In *Conference on Robot Learning*, pages 195–206, 2017.
- [26] Vijay R Konda and John N Tsitsiklis. Actor-critic algorithms. In *Advances in neural information processing systems*, pages 1008–1014, 2000.
- [27] Thanard Kurutach, Ignasi Clavera, Yan Duan, Aviv Tamar, and Pieter Abbeel. Model-ensemble trust-region policy optimization. In *International Conference on Learning Representations*, 2018.
- [28] 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.
- [29] Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. On the variance of the adaptive learning rate and beyond. *arXiv preprint arXiv:1908.03265*, 2019.
- [30] Sebastian Lunz, Ozan Öktem, and Carola-Bibiane Schönlieb. Adversarial regularizers in inverse problems. In *Advances in Neural Information Processing Systems*, pages 8507–8516, 2018.

- [31] Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.
- [32] Alberto Maria Metelli, Matteo Papini, Francesco Faccio, and Marcello Restelli. Policy optimization via importance sampling. In *Advances in Neural Information Processing Systems*, pages 5442–5454, 2018.
- [33] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937, 2016.
- [34] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pages 8024–8035, 2019.
- [35] Danil V Prokhorov and Donald C Wunsch. Adaptive critic designs. *IEEE transactions on Neural Networks*, 8(5):997–1007, 1997.
- [36] Martin L. Puterman. Markov decision processes: Discrete stochastic dynamic programming. In *Wiley Series in Probability and Statistics*, 1994.
- [37] Prajit Ramachandran, Barret Zoph, and Quoc V Le. Searching for activation functions. *arXiv preprint arXiv:1710.05941*, 2017.
- [38] Salah Rifai, Pascal Vincent, Xavier Muller, Xavier Glorot, and Yoshua Bengio. Contractive auto-encoders: Explicit invariance during feature extraction. In *ICML*, pages 833–840, 2011.
- [39] Saeed Saremi. On approximating ∇f with neural networks. *arXiv preprint arXiv:1910.12744*, 2019.
- [40] Saeed Saremi and Aapo Hyvarinen. Neural empirical bayes. *Journal of Machine Learning Research*, 20:1–23, 2019.
- [41] Jürgen Schmidhuber. Making the world differentiable: On using self-supervised fully recurrent neural networks for dynamic reinforcement learning and planning in non-stationary environments. 1990.
- [42] Jürgen Schmidhuber. Deep learning in neural networks: An overview. *Neural networks*, 61:85–117, 2015.
- [43] John Schulman, Nicolas Heess, Theophane Weber, and Pieter Abbeel. Gradient estimation using stochastic computation graphs. In *Advances in Neural Information Processing Systems*, pages 3528–3536, 2015.
- [44] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897, 2015.
- [45] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [46] Pranav Shyam, Wojciech Jaśkowski, and Faustino Gomez. Model-based active exploration. In *International Conference on Machine Learning*, pages 5779–5788, 2019.
- [47] David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. 2014.
- [48] Carl-Johann Simon-Gabriel, Yann Ollivier, Leon Bottou, Bernhard Schölkopf, and David Lopez-Paz. First-order adversarial vulnerability of neural networks and input dimension. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5809–5817, Long Beach, California, USA, 09–15 Jun 2019. PMLR.

- [49] Alice E Smith, David W Coit, Thomas Baeck, David Fogel, and Zbigniew Michalewicz. Penalty functions. *Handbook of evolutionary computation*, 97(1):C5, 1995.
- [50] Richard S Sutton. Learning to predict by the methods of temporal differences. *Machine learning*, 3(1):9–44, 1988.
- [51] Richard S Sutton. Dyna, an integrated architecture for learning, planning, and reacting. *ACM Sigart Bulletin*, 2(4):160–163, 1991.
- [52] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [53] Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pages 1057–1063, 2000.
- [54] Voot Tangkaratt, Syogo Mori, Tingting Zhao, Jun Morimoto, and Masashi Sugiyama. Model-based policy gradients with parameter-based exploration by least-squares conditional density estimation. *Neural networks*, 57:128–140, 2014.
- [55] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012.
- [56] John N Tsitsiklis and Benjamin Van Roy. Analysis of temporal-difference learning with function approximation. In *Advances in neural information processing systems*, pages 1075–1081, 1997.
- [57] Hado P van Hasselt, Matteo Hessel, and John Aslanides. When to use parametric models in reinforcement learning? In *Advances in Neural Information Processing Systems*, pages 14322–14333, 2019.
- [58] Xin Wang and Thomas G Dietterich. Model-based policy gradient reinforcement learning. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pages 776–783, 2003.
- [59] Théophane Weber, Nicolas Heess, Lars Buesing, and David Silver. Credit assignment techniques in stochastic computation graphs. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2650–2660, 2019.
- [60] Paul Werbos. Advanced forecasting methods for global crisis warning and models of intelligence. *General Systems Yearbook*, 22, 01 1977.
- [61] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.

A Proof of Proposition 3.1

The proof follows directly from the Deterministic Policy Gradient Theorem, Therefore, the Proposition inherits all of its smoothness assumptions about the Markov Decision Process [47].

Proposition 3.1. *Let Π be a parametric space of L_π -Lipschitz continuous differentiable deterministic policies, \mathcal{Q} a space of approximate value functions and $\|\cdot\|$ any p -norm. Given $\pi \in \Pi$ and $\hat{Q} \in \mathcal{Q}$, the norm of the difference between the true policy gradient $\nabla_{\theta} J(\theta)$ and its approximation $\hat{\nabla}_{\theta} J(\theta)$, which uses \hat{Q} , can be upper bounded as:*

$$\|\nabla_{\theta} J(\theta) - \hat{\nabla}_{\theta} J(\theta)\| \leq \frac{L_\pi}{1-\gamma} \mathbb{E}_{s \sim d_\mu^\pi} \left\| \nabla_a \delta^{\pi, \hat{Q}}(s, a) \Big|_{a=\pi(s)} \right\|.$$

Proof.

$$\|\nabla_{\theta} J(\theta) - \hat{\nabla}_{\theta} J(\theta)\| = \frac{1}{1-\gamma} \left\| \int_{\mathcal{S}} d_\mu^\pi(s) \left(\nabla_a Q^\pi(s, a) \Big|_{a=\pi(s)} - \nabla_a \hat{Q}(s, a) \Big|_{a=\pi(s)} \right) \nabla_{\theta} \pi(s) ds \right\| \quad (10)$$

$$= \frac{1}{1-\gamma} \left\| \int_{\mathcal{S}} d_\mu^\pi(s) \nabla_a \delta^{\pi, \hat{Q}}(s, a) \Big|_{a=\pi(s)} \nabla_{\theta} \pi(s) ds \right\| \quad (11)$$

$$\leq \frac{1}{1-\gamma} \int_{\mathcal{S}} d_\mu^\pi(s) \left\| \nabla_a \delta^{\pi, \hat{Q}}(s, a) \Big|_{a=\pi(s)} \right\| \cdot \|\nabla_{\theta} \pi(s)\| ds$$

$$\leq \frac{L_\pi}{1-\gamma} \int_{\mathcal{S}} d_\mu^\pi(s) \left\| \nabla_a \delta^{\pi, \hat{Q}}(s, a) \Big|_{a=\pi(s)} \right\| ds. \quad (12)$$

□

Equation 10 follows from the Deterministic Policy Gradient Theorem. To obtain Equation 11, we exploit the definition of $\delta^{\pi, \hat{Q}}$ and linearity of differentiation. Finally, in Equation 12, we use the Lipschitz policy assumption.

B Additional Experiments

B.1 Unconstrained Action-Value Gradient learning

Proposition 3.1 directly encourages training the critic by minimizing the bound on the error of the policy gradient, i.e., the norm of the action-gradient of the policy evaluation error.

However, we found a direct optimization of this bound, by means of the TD-error, difficult in the context of Dyna-like algorithms. We analyze this behavior in the Pendulum-v0 environment [6], instantiating a version of MAGE based on DDPG [28] (MAGE-DDPG). To understand the learning dynamics of the action-value gradients in a way that is not affected by the model bias, we employ the differentiable version of the real environment dynamics and test MAGE without the TD-error regularization (i.e., with $\lambda = 0$). Therefore, at each step, Q_ω is improved by minimizing the norm of $\hat{\delta}$ computed on transitions whose next state is sampled from p . Unfortunately, no useful learning can be achieved in this setting: a degenerate solution consisting of \hat{Q} such that $\left\| \nabla_a \hat{Q}(s, a) \right\| \approx 0, \forall s \in \mathcal{S}, \forall a \in \mathcal{A}$ is rapidly reached, as shown in Figure 5. We employ exactly the settings and hyperparameters that are successfully employed in the full version of MAGE.

We believe that understanding whether, or under which circumstances, the direct minimization of the bound in Proposition 3.1 is possible is an interesting open question.

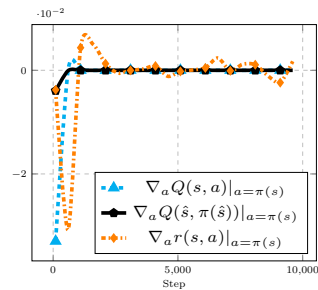


Figure 5: Action-gradients

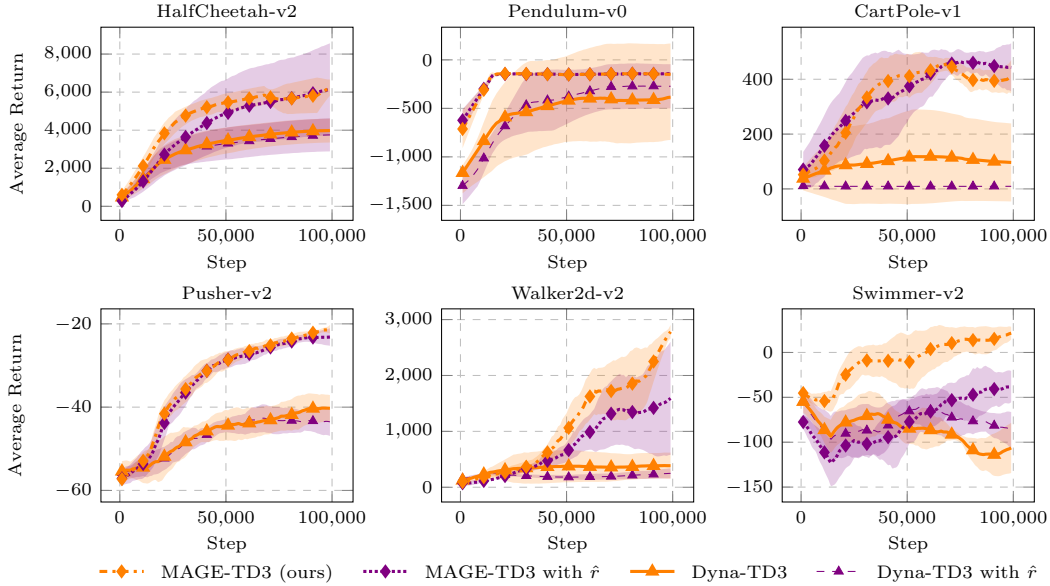


Figure 6: Performance in terms of average return of MAGE-TD3 and Dyna-TD3 with and without the use of an estimated reward function \hat{r} (5 runs, 95% c.i.).

B.2 MAGE with Trained Reward Function

As discussed in Section 4, MAGE is able to achieve good performance even with an estimated reward function. We report in Figure 6 the full results of this experiment on all the considered environments. For reference, we test MAGE and Dyna-TD3 as well as their versions in which the ground-truth reward function is substituted with one trained on the experience replay data using the MSE loss.

The results indicate that learning the reward function when it is not directly accessible does not produce any catastrophic harm to the performance of the algorithm. Therefore, our approach remains competitive even when the assumption of a know differentiable reward function is not satisfied.

B.3 Importance of model capacity

The quality of the learned model is of paramount importance for most MBRL algorithms, whose performance, generally, deteriorates when the model is not enough expressive for a given task. Thus, we performed an additional experiment to investigate how the performance of MAGE (compared to the Dyna-TD3 baseline) is affected by the use of less powerful models. We evaluated two versions of the model with reduced capacity: (i) only 2 members in the ensemble, 4 hidden layers and 256 units per layer (*-small suffix*); (ii) no ensemble (a single model), 2 hidden layers, 256 units per layer (*-smaller suffix*). Recall that the original setting involves a more powerful model with 8 members in the ensemble, 4 hidden layers, 512 units per layer (no suffix). The results on the Pusher-v2 environment, reported in Figure 7, show that MAGE is robust to the presence of a misspecified model: while a simpler but still quite capable model does no harm to MAGE, a significantly smaller model has a reasonable impact on the obtained average return.

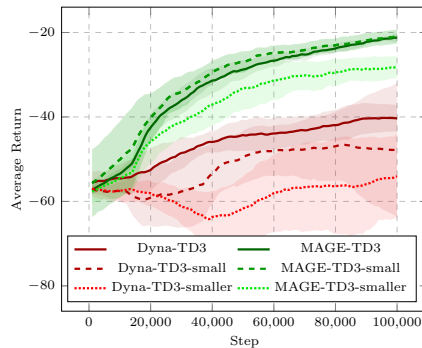


Figure 7: Pusher-v2 (5 runs, 95% c.i.).

B.4 Importance of λ

Our practical solution to viably minimize the norm of the action-gradient of the TD-error involves a constrained optimization problem, that limits the magnitude of the traditional TD-error. We approximately solve this problem by transforming it into an unconstrained one, introducing a new hyperparameter λ . λ can be seen as a weight that is given to the traditional TD-error, assigning more or less importance to it compared to the error on the action-gradient. In the main experiment shown in the paper, we used $\lambda = 0.2$, which was chosen arbitrarily. How sensitive is MAGE to this parameter?

To study that, we carried out an experiment on the environment HalfCheetah-v2, by testing the TD3-based version of MAGE using four different values of λ . The results are shown in Figure 8, and demonstrate that, regardless of the value of λ , MAGE is significantly better than the baseline Dyna-TD3. MAGE is therefore robust to the choice of this hyperparameter. Notice also that the $\lambda = 0.2$ we used is probably not optimal for some environments: thus, the absolute returns obtained by MAGE could be improved for particular tasks if a different hyperparameter is used for each of them, which we leave for a future work. Nonetheless, we decided to report in Figure 2 results for a fixed value of λ across all the environments to show the robustness and ease of use of MAGE.

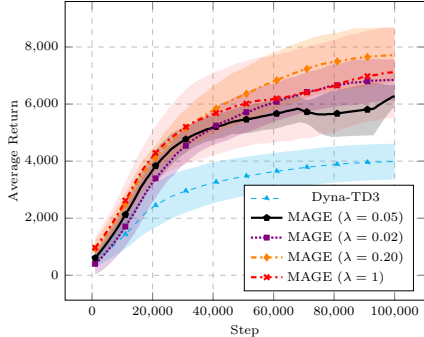


Figure 8: Median return of MAGE for different $\lambda \in [0, 1]$ (5 runs, 95% c.i.).

B.5 Asymptotic Performance

A particularly important question concerning the performance of a model-based reinforcement learning algorithm is whether it matches the one of model-free baselines. We answer this question by running MAGE-TD3 and Dyna-TD3 until convergence on the HalfCheetah-v2 environment and replicating the evaluation procedure used to obtain the asymptotic performance of an algorithm in the original TD3 paper [18]. This amounts to repeating a number of trials (we use 5 of them for both the model-based algorithms) and considering the maximum average return over them.

Table 1: Maximum Average Return on HalfCheetah-v2 over different trials for three versions of TD3. Best performance is in bold.

Algorithm	Maximum Average Return
MAGE-TD3 ($3 \cdot 10^6$ steps, 5 trials)	9660.79 ± 2821.31
Dyna-TD3 ($3 \cdot 10^6$ steps, 5 trials)	8372.76 ± 1859.73
TD3 ($1 \cdot 10^7$ steps, 10 trials)	9636.95 ± 859.07

Results are reported in Table 1. Two conclusions can be drawn from them. The first one is that MAGE not only shows superior sample-efficiency compared to its model-free counterpart, but also matches its asymptotic performance in a smaller number of steps (less than one third). Secondly, the inferior performance of Dyna-TD3 compared to model-free TD3 once again reinforces the evidence that any simple introduction of a dynamics model into a model-free algorithm does not guarantee an improvement, when measuring performance using commonly employed metrics.

C Action-Gradient of the TD-error

In this section, we present some additional information about the computation of the action-gradient of the TD-error, carried out during the critic learning step of MAGE. To implement MAGE, we employed PyTorch [34] and its automatic differentiation tools in order to compute the second-order gradient required by our method. In this way, we did not need to explicitly derive a closed form expression for a given model class or neural network architecture. Nonetheless, we report here the

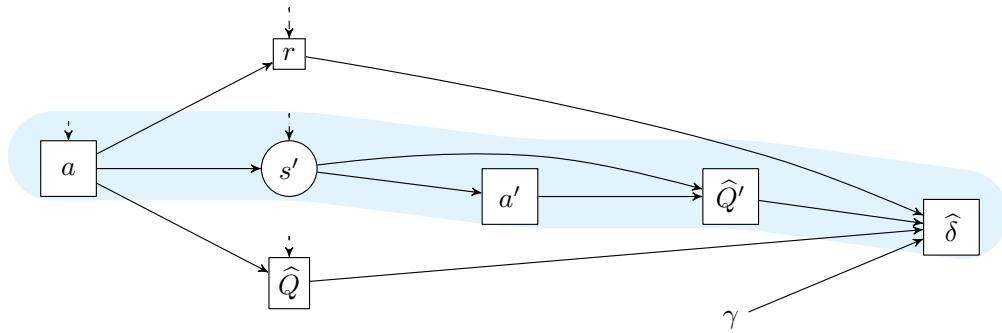


Figure 9: Alternative view of the computational graph constructed during the computation of the TD-error $\hat{\delta}$, following the notation from [43]. Round nodes represent stochastic variables, squares represent deterministic variables. Nodes with incoming dashed edges also depend on the state s .

general expression for the action-gradient of the TD-error:

$$\frac{\partial \hat{\delta}^{\pi, \hat{Q}}(s, a, s')}{\partial a} = \frac{\partial r(s, a)}{\partial a} + \gamma \frac{\partial \hat{p}(s'|s, a)}{\partial a} \left(\frac{\partial \hat{Q}(s', \pi(s'))}{\partial s'} + \frac{\partial \pi(s')}{\partial s'} \frac{\partial \hat{Q}(s', a')}{\partial a'} \right) - \frac{\partial \hat{Q}(s, a)}{\partial a}. \quad (13)$$

In MAGE, we employ a Gaussian stochastic model \hat{p} : therefore, its action-gradient $\frac{\partial \hat{p}(s'|s, a)}{\partial a}$ can be obtained by reparameterizing this distribution using randomly drawn unit Gaussian noise together with the learned mean and standard deviations. In our experiments, we only deal with continuous state and action spaces; however, by leveraging appropriate approximations (e.g., concrete distributions [31]), similar techniques can be employed also in the case of a discrete state space \mathcal{S} .

To further visualize the constructed computational graph, it is possible to employ a different view, inspired recent work on stochastic computational graphs [43], w.r.t. the one leveraged in Figure 1 (see Figure 9). In our case, the only possibly stochastic entity is the approximate model.

D Experimental details

D.1 Instantiating MAGE

We presented in Algorithm 1 a generic version with MAGE, whose structure can be adapted to many model-free actor-critic algorithms. In most of our experiments, we use TD3 [18] as a reference algorithm, due to its stability and performance, giving birth to MAGE-TD3. In Algorithm 2, we report pseudocode for this version of our method. Unfortunately, while the use of the model is unchanged w.r.t. the abstract version, the addition of a second value function implies the computational overhead of using second-order differentiation twice.

D.2 Hyperparameters

We employ 1000 (100 for the Pendulum and Cartpole environments) warmup steps of interaction with the environment before starting to update the critic and the actor. We use an ensemble of 8 neural network as approximate dynamics models, that learn both mean and standard deviation of a Gaussian distribution, of 4 hidden layers of 512 neurons (2 layers with 128 units for the Pendulum and Cartpole environments) with swish [37] activation function. They are trained by maximum likelihood, minimizing the loss function, after every 25 steps of interaction with the environment, on 120 batches of 256 samples. We employ multi-layer perceptrons also for the actor (2 layers, 128 neurons each for the Pendulum and CartPole environments and 284 for all the others) and the critic (2 layers, 384 neurons each). Model, actor and critic are trained with the RAdam optimizer [29], with learning rates of 0.0001 and default parameters, and a weight decay of 0.0001 for the approximate dynamics model. For MAGE-TD3, we employ $\lambda = 0.2$ for the experiment showed in Figure 2 and $\lambda = 0.05$ for the other experiments. We update the critic and the actor by extracting 1024 (512 for the

Algorithm 2 Model-based Action-Gradient-Estimator TD3 (MAGE-TD3)

Input: Initial buffer \mathcal{B} , parameters $\{\omega, \phi_1, \phi_2, \theta\}$, target parameters $\{\bar{\phi}_1 = \phi_1, \bar{\phi}_2 = \phi_2, \bar{\theta} = \theta\}$

for each iteration **do**
Collect transition (s, a, s') acting according to exploratory policy $\pi_\epsilon(s) = \pi_\theta(s) + \epsilon, \epsilon \sim \mathcal{N}(0, \sigma)$
 $\mathcal{B} \leftarrow \mathcal{B} \cup \{(s, a, s')\}$
for each model learning step **do**
 $\omega \leftarrow \omega - \alpha_p \nabla_\omega \ell(s, a, s'; \omega), \quad (s, a, s') \sim \mathcal{B}$
end for
for each policy optimization step **do**
Extract state s after sampling $(s, \cdot, \cdot) \sim \mathcal{B}$
 $\hat{y} \leftarrow r(s, \pi_\epsilon(s)) + \gamma \min_{i=1,2} Q_{\bar{\phi}_i}(\hat{s}, \pi_{\bar{\theta}}(\hat{s})), \quad \hat{s} \sim p_\omega(\cdot|s, \pi_\theta(s)), \epsilon \sim \text{clip}(\mathcal{N}(0, \bar{\sigma}))$
for $i \in \{1, 2\}$ **do**
 $\hat{\delta}(s, a, \hat{s}; \phi_i) \leftarrow \hat{y} - Q_{\phi_i}(s, a), \quad a = \pi_\theta(s)$
 $\phi_i \leftarrow \phi - \alpha_Q \nabla_{\phi_i} \left(\left\| \nabla_a \hat{\delta}(s, a, \hat{s}; \phi_i) \Big|_{a=\pi_\theta(s)} \right\| + \lambda \left| \hat{\delta}(s, a, \hat{s}; \phi_i) \right| \right)$
 $\bar{\phi}_i \leftarrow \tau \phi_i + (1 - \tau) \bar{\phi}_i$
end for
if $t \bmod d = 0$ **then**
 $\theta \leftarrow \theta + \alpha_\pi \nabla_\theta \min_{i=1,2} Q_{\phi_i}(s, \pi_\theta(s))$
 $\bar{\theta} \leftarrow \tau \theta + (1 - \tau) \bar{\theta}$
end if
end for
end for

Pendulum and CartPole environment) states from the buffer of collected transitions, then sampling from the ensemble by first randomly selecting one of the members and then sampling an estimated difference between current and next state. The critic is trained by employing an Huber loss.

In MAGE-TD3, we employ the suggested hyperparameters of TD3: an action noise of 0.1, a target noise of 0.2, noise clipping to 0.5 and a delay in the policy updates of 2. During training, the actions that the actor executes in the environment are perturbed by Gaussian noise $\epsilon \sim \mathcal{N}(0, 0.1)$. We obtain the target networks for both actor and critic by Polyak averaging with decay $\tau = 0.995$.

For the reward estimation experiments, we employ a neural network with 3 hidden layers of 256 units (1 hidden layer with 128 units for the Pendulum and Cartpole environments) and swish activations. We employ a discount factor of $\gamma = 0.99$.

For our experiment on the evaluation of gradients, we initially collect 200 transitions, then simply run the algorithms with standard settings but without any update of the actor. Every 10 steps, we collect 10 trajectories in the environment and average the error over them. We compute the ground-truth $\nabla_a G(s, a)|_{a=\pi(s)}$, with $G(s, a) = \sum_{t=0}^{H-1} \gamma^t r(s_t, a_t)|_{s=s_0, a=a_0}$ being the empirical return, by automatic differentiation, leveraging the differentiable oracle model. We then average, the resulting discounted error:

$$L(Q^\pi, \hat{Q}) = \frac{1}{H} \sum_{t=0}^{H-1} \gamma^t \|\nabla_a G(s_t, a_t) - \nabla_a \hat{Q}(s_t, a_t)\|_1. \quad (14)$$

We average this value across the 10 different trajectories.

Across all the experiments, despite the formulation we used throughout the paper, we employ a reward $r(s, a, s')$, which is thus also a function of the next state. Formally, it is possible to interpret the state-action reward we use throughout the paper as $r(s, a) = \mathbb{E}_{s' \sim p(\cdot|s, a)} [r(s, a, s')]$. For generating the performance plots, we evaluate, after every 1000 steps of environment interaction, the actor for 10 episodes and average the result. To improve presentation, we then uniformly smooth the resulting curves with a window size of 25.