

# Deterministic Value-Policy Gradients

Qingpeng Cai,<sup>\*1</sup> Ling Pan,<sup>\*2</sup> Pingzhong Tang<sup>2</sup>

<sup>1</sup>Alibaba Group

<sup>2</sup>IIS, Tsinghua University

qingpeng.cqp@alibaba-inc.com, pl17@mails.tsinghua.edu.cn, kenshin@tsinghua.edu.cn

## Abstract

Reinforcement learning algorithms such as the deep deterministic policy gradient algorithm (DDPG) has been widely used in continuous control tasks. However, the model-free DDPG algorithm suffers from high sample complexity. In this paper we consider the deterministic value gradients to improve the sample efficiency of deep reinforcement learning algorithms. Previous works consider deterministic value gradients with the finite horizon, but it is too myopic compared with infinite horizon. We firstly give a theoretical guarantee of the existence of the value gradients in this infinite setting. Based on this theoretical guarantee, we propose a class of the deterministic value gradient algorithm (DVG) with infinite horizon, and different rollout steps of the analytical gradients by the learned model trade off between the variance of the value gradients and the model bias. Furthermore, to better combine the model-based deterministic value gradient estimators with the model-free deterministic policy gradient estimator, we propose the deterministic value-policy gradient (DVPG) algorithm. We finally conduct extensive experiments comparing DVPG with state-of-the-art methods on several standard continuous control benchmarks. Results demonstrate that DVPG substantially outperforms other baselines.

## Introduction

(Silver et al. 2014) propose the deterministic policy gradient (DPG) algorithm that aims to find an optimal deterministic policy that maximizes the expected long-term reward, which lowers the variance when estimating the policy gradient, compared to stochastic policies (Sutton et al. 2000). (Lillicrap et al. 2015) further combine deep neural networks with DPG to improve the modeling capacity, and propose the deep deterministic policy gradient (DDPG) algorithm. It is recognized that DDPG has been successful in robotic control tasks (Gu et al. 2016) and dynamic mechanism design (Tang 2017; Cai et al. 2018). Despite the effectiveness of DDPG in these tasks, it suffers from the high sample complexity problem (Schulman et al. 2015).

Deterministic value gradient methods (Werbos 1990; Nguyen and Widrow 1990; Jordan and Rumelhart 1992;

Fairbank 2008) compute the policy gradient through back propagation of the reward along a trajectory predicted by the learned model, which enables better sample efficiency. However, to the best of our knowledge, existing works of deterministic value gradient methods merely focus on finite horizon, which are too myopic and can lead to large bias. Stochastic value gradient (SVG) methods (Heess et al. 2015) use the re-parameterization technique to optimize the stochastic policies. Among the class of SVG algorithms, although SVG(1) studies infinite-horizon problems, it only uses one-step rollout, which limits its efficiency. Also, it suffers from the high variance due to the importance sampling ratio and the randomness of the policy.

In this paper, we study the setting with infinite horizon, where both state transitions and policies are deterministic. (Heess et al. 2015) gives recursive Bellman gradient equations of deterministic value gradients, but the gradient lacks of theoretical guarantee as the DPG theorem does not hold in this deterministic transition case. We prove that the gradient indeed exists for a certain set of discount factors. We then derive a closed form of the value gradients.

However, the estimation of the deterministic value gradients is much more challenging. The difficulty of the computation of the gradient mainly comes from the dependency of the gradient of the value function over the state. Such computation may involve infinite times of the product of the gradient of the transition function and is hard to converge. Thus, applying the Bellman gradient equation recursively may incur high instability.

To overcome these challenges, we use model-based approaches to predict the reward and transition function. Based on the theoretical guarantee of the closed form of the value gradients in the setting, we propose a class of deterministic value gradients  $DVG(k)$  with infinite horizon, where  $k$  denotes the number of rollout steps. For each choice of  $k$ , we use the rewards predicted by the model and the action-value at  $k + 1$  step to estimate of the value gradients over the state, in order to reduce the instability of the gradient of the value function over the state. Different number of rollout steps maintains a trade-off between the accumulated model bias and the variance of the gradient over the state. The deterministic policy gradient estimator can be viewed as a special case

<sup>\*</sup>The first two authors contributed equally to this work.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

of this class, i.e., it never use the model to estimate the value gradients, and we refer it to DVG(0).

As the model-based approaches are more sample efficient than model-free algorithms (Li and Todorov 2004; Levine and Koltun 2013), and the model-based deterministic value gradients may incur model bias (Wahlström, Schön, and Deisenroth 2015), we consider an essential question: *How to combines the model-based gradients and the model-free gradients efficiently?*

We propose a temporal difference method to ensemble gradients with different rollout steps. The intuition is to ensemble different gradient estimators with geometric decaying weights. Based on this estimator, we propose the deterministic value-policy gradient (DVPG) algorithm. The algorithm updates the policy by stochastic gradient ascent with the ensembled value gradients of the policy, and the weight maintains a trade-off between sample efficiency and performance.

To sum up, the main contribution of the paper is as follows:

- First of all, we provide a theoretical guarantee for the existence of the deterministic value gradients in settings with infinite horizon.
- Secondly, we propose a novel algorithm that ensembles the deterministic value gradients and the deterministic policy gradients, called deterministic value-policy gradient (DVPG), which effectively combines the model-free and model-based methods. DVPG reduces sample complexity, enables faster convergence and performance improvement.
- Finally, we conduct extensive experiments on standard benchmarks comparing with DDPG, DDPG with model-based rollouts, the stochastic value gradient algorithm, SVG(1) and state-of-the-art stochastic policy gradient methods. Results confirm that DVPG significantly outperforms other algorithms in terms of both sample efficiency and performance.

## Related Work

Model-based algorithms has been widely studied (Moldovan et al. 2015; Montgomery and Levine 2016; Ha and Schmidhuber 2018; Hafner et al. 2018; Chua et al. 2018) in recent years. Model-based methods allows for more efficient computations and faster convergence than model-free methods (Wang and Dietterich 2003; Li and Todorov 2004; Levine and Koltun 2013; Watter et al. 2015).

There are two classes of model-based methods, one is to use learned model to do imagination rollouts to accelerate the learning. (Gu et al. 2016; Kurutach et al. 2018) generate synthetic samples by the learned model. PILCO (Deisenroth and Rasmussen 2011) learns the transition model by Gaussian processes and applies policy improvement on analytic policy gradients. The other is to use learned model to get better estimates of action-value functions. The value prediction network (VPN) (Oh, Singh, and Lee 2017) uses the learned transition model to get a better target estimate. (Feinberg et al. 2018; Buckman et al. 2018) combines different model-based value expansion functions by TD( $k$ ) trick or stochastic distributions to improve the estimator of the action-value function.

Different from previous model-based methods, we present a temporal difference method that ensembles model-based deterministic value gradients and model-free policy gradients. Our technique can be combined with both the imagination rollout technique and the model-based value expansion technique.

## Preliminaries

A Markov decision process (MDP) is a tuple  $(S, A, p, r, \gamma, p_0)$ , where  $S$  and  $A$  denote the set of states and actions respectively.  $p(s_{t+1}|s_t, a_t)$  represents the conditional density from state  $s_t$  to state  $s_{t+1}$  under action  $a_t$ . The density of the initial state distribution is denoted by  $p_0(s)$ . At each time step  $t$ , the agent interacts with the environment with a deterministic policy  $\mu_\theta$ . We use  $r(s_t, a_t)$  to represent the immediate reward, contributing to the discounted overall rewards from state  $s_0$  following  $\mu_\theta$ , denoted by  $J(\mu_\theta) = \mathbb{E}[\sum_{k=0}^{\infty} \gamma^k r(a_k, s_k) | \mu_\theta, s_0]$ . Here,  $\gamma \in [0, 1)$  is the discount factor. The Q-function of state  $s_t$  and action  $a_t$  under policy  $\mu_\theta$  is denoted by  $Q^{\mu_\theta}(s_t, a_t) = \mathbb{E}[\sum_{k=t}^{\infty} \gamma^{k-t} r(a_k, s_k) | \mu_\theta, s_t, a_t]$ . The corresponding value function of state  $s_t$  under policy  $\mu_\theta$  is denoted by  $V^{\mu_\theta}(s_t) = Q^{\mu_\theta}(s_t, \mu_\theta(s_t))$ . We denote the density at state  $s'$  after  $t$  time steps from state  $s$  following the policy  $\mu_\theta$  by  $p(s, s', t, \mu_\theta)$ . We denote the discounted state distribution by  $\rho^{\mu_\theta}(s') = \int_S \sum_{t=1}^{\infty} \gamma^{t-1} p_0(s) p(s, s', t, \mu_\theta) ds$ . The agent aims to find an optimal policy that maximizes  $J(\mu_\theta)$ .

## Deterministic Value Gradients

In this section, we study a setting of infinite horizon with deterministic state transition, which poses challenges for the existence of deterministic value gradients. We first prove that under proper condition, the deterministic value gradient does exist. Based on the theoretical guarantee, we then propose a class of practical algorithms by rolling out different number of steps. Finally, we discuss the difference and connection between our proposed algorithms and existing works.

Deterministic Policy Gradient (DPG) Theorem (Silver et al. 2014), proves the existence of the deterministic policy gradient for MDP that satisfies the regular condition, which requires the probability density of the next state  $p(s'|s, a)$  to be differentiable in  $a$ . In the proof of the DPG theorem, the existence of the gradient of the value function is firstly proven, i.e.,

$$\nabla_\theta V^{\mu_\theta}(s) = \int_S \sum_{t=0}^{\infty} \gamma^t p(s, s', t, \mu_\theta) \nabla_\theta \mu_\theta(s') \nabla_{a'} Q^{\mu_\theta}(s', a')|_{a'=\mu_\theta(s')} ds', \quad (1)$$

then the gradient of the long-terms rewards exists. Without this condition, the arguments in the proof of the DPG theorem do not work<sup>1</sup>, and poses challenges for cases where the differentiability is not satisfied. Note this condition does not hold in any case with deterministic transitions. Therefore, one must need a new theoretical guarantee to determine the

<sup>1</sup>Please refer to <http://proceedings.mlr.press/v32/silver14-supplement.pdf>

existence of the gradient of  $V^{\mu_\theta}(s)$  over  $\theta$  in deterministic state transition cases.

### Deterministic Value Gradient Theorem

We now analyze the gradient of a deterministic policy. Denote  $T(s, a)$  the next state given current state  $s$  and action  $a$ . Without loss of generality, we assume that the transition function  $T$  is continuous, differentiable in  $s$  and  $a$  and is bounded. Note that the regular condition is not equivalent to this assumption. Consider a simple example that a transition  $T(s, a) = s + a$ , then the gradient of  $p(s'|s, a)$  over  $a$  is infinite or does not exist. However, the gradient of  $T(s, a)$  over  $a$  exists. By definition,

$$\begin{aligned}\nabla_\theta V^{\mu_\theta}(s) &= \nabla_\theta \left( r(s, \mu_\theta(s)) + \gamma V^{\mu_\theta}(s') \right) \Big|_{s'=T(s, \mu_\theta(s))} \\ &= \nabla_\theta r(s, \mu_\theta(s)) + \gamma \nabla_\theta V^{\mu_\theta}(s') \Big|_{s'=T(s, \mu_\theta(s))} \\ &\quad + \gamma \nabla_\theta T(s, \mu_\theta(s)) \nabla_{s'} V^{\mu_\theta}(s').\end{aligned}$$

Therefore, the key of the existence (estimation) of the gradient of  $V^{\mu_\theta}(s)$  over  $\theta$  is the existence (estimation) of  $\nabla_{s'} V^{\mu_\theta}(s')$ . In Theorem 1, we give a sufficient condition of the existence of  $\nabla_{s'} V^{\mu_\theta}(s')$ .

**Theorem 1** For any policy  $\mu_\theta$ , the gradient of the value function over the state,  $\nabla_{s'} V^{\mu_\theta}(s')$ , exists with two assumptions:

- A.1: The set of states that the policy visits starting from any initial state  $s$  is finite.
- A.2: For any initial state  $s$ , by Assumption A.1, we get that there is a periodic loop of visited states. Let  $(s_0, s_1, \dots, s_k)$  denote the loop<sup>2</sup>, and  $A(s) = \gamma^{k+1} \prod_{i=0}^k \nabla_{s_i} T(s_i, \mu_\theta(s_i))$ , the power sum of  $A(s)$ ,  $\sum_{m=0}^{\infty} A^m(s)$  converges.

**Proof 1** By definition,

$$V^{\mu_\theta}(s) = r(s, \mu_\theta(s)) + \gamma V^{\mu_\theta}(s') \Big|_{s'=T(s, \mu_\theta(s))}. \quad (2)$$

Taking the gradient of Eq. (2), we obtain

$$\begin{aligned}\nabla_{s'} V^{\mu_\theta}(s) &= \nabla_{s'} r(s, \mu_\theta(s)) \\ &\quad + \gamma \nabla_{s'} T(s, \mu_\theta(s)) \nabla_{s''} V^{\mu_\theta}(s'') \Big|_{s''=T(s, \mu_\theta(s))}.\end{aligned} \quad (3)$$

Unrolling Eq. (3) with infinite steps, we get

$$\nabla_{s'} V^{\mu_\theta}(s) = \sum_{t=0}^{\infty} \gamma^t g(s, t, \mu_\theta) \nabla_{s_t} r(s_t, \mu_\theta(s_t)), \quad (4)$$

where  $g(s, t, \mu_\theta) = \prod_{i=0}^{t-1} \nabla_{s_i} T(s_i, \mu_\theta(s_i))$ ,  $s_0 = s$  and  $s_i$  is the state after  $i$  steps following policy  $\mu_\theta$ .

With the assumption A.1, we rewrite (4) by the indicator function  $I(s, s', t, \mu_\theta)$  that indicates whether  $s'$  is obtained after  $t$  steps from the initial state  $s$  following the policy  $\mu_\theta$ :

$$\begin{aligned}\nabla_{s'} V^{\mu_\theta}(s) &= \sum_{t=0}^{\infty} \sum_{s' \in B(s, \theta)} \gamma^t g(s, t, \mu_\theta) I(s, s', t, \mu_\theta) \\ &\quad \nabla_{s'} r(s', \mu_\theta(s')), \end{aligned} \quad (5)$$

<sup>2</sup>Note that  $s_0$  is not equal to  $s_k$ .

Where  $B(s, \theta)$  is the set of states the policy visits from  $s$ .

We now prove that for any  $\mu_\theta, s, s'$ , the infinite sum of gradients,  $\sum_{t=0}^{\infty} \gamma^t g(s, t, \mu_\theta) I(s, s', t, \mu_\theta)$  converges.

For each state  $s'$ , there are three cases during the process from the initial state  $s$  with infinite steps:

1. Never visited:  $\sum_{t=0}^{\infty} \gamma^t g(s, t, \mu_\theta) I(s, s', t, \mu_\theta) = 0$ .
2. Visited once: Let  $t_{s'}$  denote the number of steps that it takes to reach the state  $s'$ , then  $\sum_{t=0}^{\infty} \gamma^t g(s, t, \mu_\theta) I(s, s', t, \mu_\theta) = \gamma^{t_{s'}} g(s, t_{s'}, \mu_\theta)$ .
3. Visited infinite times: Let  $t_1$  denote the number of steps it takes to reach  $s'$  for the first time. The state  $s'$  will be re-visited every  $k$  steps after the previous visit. By definition,

$$\begin{aligned}&\sum_{t=0}^{\infty} \gamma^t g(s, t, \mu_\theta) I(s, s', t, \mu_\theta) \\ &= \sum_{a=0}^{\infty} \gamma^{t_1+a} g(s, t_1+a, \mu_\theta) A^a(s).\end{aligned} \quad (6)$$

By the assumption A.2 we get (6) converges.

By exchanging the order of the limit and the summation,

$$\begin{aligned}\nabla_{s'} V^{\mu_\theta}(s) &= \sum_{s' \in B(s, \theta)} \sum_{t=0}^{\infty} \gamma^t g(s, t, \mu_\theta) I(s, s', t, \mu_\theta) \\ &\quad \nabla_{s'} r(s', \mu_\theta(s')).\end{aligned} \quad (7)$$

Assumption A.1 guarantees the existence of the stationary distribution of states theoretically. Actually, it holds on most continuous tasks, e.g., InvertedPendulum-v2 in MuJoCo. We directly test a deterministic policy with a 2-layer fully connected network on this environment with 10,000 episodes<sup>3</sup>, and we count the number that each state is visited. After projecting the data into 2D space by t-SNE (Maaten and Hinton 2008), we obtain the state visitation density contour as shown in Figure 1. We have two interesting findings: (1) The set of states visited by the policy is finite. (2) Many states are visited for multiple times, which justifies Assumption A.1.

By the analysis of Assumption A.2, we get that for any policy and state, there exists a set of discount factors such that the gradient of the value function over the state exists, as illustrated in Corollary 1.

**Corollary 1** For any policy  $\mu_\theta$  and any initial state  $s$ , let  $(s_0, s_1, \dots, s_k)$  denote the loop of states following the policy and the state,  $C(s, \mu_\theta, k) = \prod_{i=0}^k \nabla_{s_i} T(s_i, \mu_\theta(s_i))$ , the gradient of the value function over the state,  $\nabla_{s'} V^{\mu_\theta}(s)$  exists if  $\gamma^{k+1} \max \{ \|C(s, \mu_\theta, k)\|_\infty, \|C(s, \mu_\theta, k)\|_1 \} < 1$ .

In Theorem 2, we show that the deterministic value gradients exist and obtain the closed form based on the analysis in Theorem 1.

**Theorem 2 (Deterministic Value Gradient Theorem)** For any policy  $\mu_\theta$  and MDP with deterministic state transitions,

<sup>3</sup>We test different weights, the observation of finite visited states set is very common among different weights.




Figure 1: State visitation density contour on InvertedPendulum-v2.

if assumptions A.1 and A.2 hold, the value gradients exist, and

$$\begin{aligned} \nabla_{\theta} V^{\mu_{\theta}}(s) = & \sum_{s' \in B(s, \theta)} \rho^{\mu_{\theta}}(s, s') \nabla_{\theta} \mu_{\theta}(s') (\nabla_{a'} r(s', a') + \\ & \gamma \nabla_{a'} T(s', a') \nabla_{s''} V^{\mu_{\theta}}(s'')|_{s''=T(s', a')}, \end{aligned}$$

where  $a'$  is the action the policy takes at state  $s'$ ,  $\rho^{\mu_{\theta}}(s, s')$  is the discounted state distribution starting from the state  $s$  and the policy, and is defined as  $\rho^{\mu_{\theta}}(s, s') = \sum_{t=1}^{\infty} \gamma^{t-1} I(s, s', t, \mu_{\theta})$ .

### Deterministic Value Gradient Algorithm

The value gradient methods estimate the gradient of value function recursively (Fairbank and Alonso 2012):

$$\begin{aligned} \nabla_{\theta} V^{\mu_{\theta}}(s) = & \nabla_{\theta} r(s, \mu_{\theta}(s)) + \gamma \nabla_{\theta} T(s, \mu_{\theta}(s)) \nabla_{s'} V^{\mu_{\theta}}(s') \\ & + \gamma \nabla_{\theta} V^{\mu_{\theta}}(s') \end{aligned} \quad (8)$$

$$\begin{aligned} \nabla_s V^{\mu_{\theta}}(s) = & \nabla_s r(s, \mu_{\theta}(s)) + \gamma \nabla_s T(s, \mu_{\theta}(s)) \\ & \nabla_{s'} V^{\mu_{\theta}}(s')|_{s'=T(s, \mu_{\theta}(s))}. \end{aligned} \quad (9)$$

In fact, there are two kinds of approaches for estimating the gradient of the value function over the state, i.e., infinite and finite. On the one hand, directly estimating the gradient of the value function over the state recursively by Eq. (9) for infinite times is slow to converge. On the other hand, estimating the gradient by finite horizon like traditional value gradient methods (Werbos 1990; Nguyen and Widrow 1990; Heess et al. 2015) may cause large bias of the gradient.

We set out to estimate the action-value function denoted by  $Q^w(s, a)$  with parameter  $w$ , and replace  $\nabla_s V^{\mu_{\theta}}(s)$  by  $\nabla_s Q^w(s, \mu_{\theta}(s))$  in Eq. (8). In this way, we can directly obtain a 1-step estimator of the value gradients,

$$\begin{aligned} G_1(\mu_{\theta}, s) = & \nabla_{\theta} r(s, \mu_{\theta}(s)) + \gamma \nabla_{\theta} T(s, \mu_{\theta}(s)) \\ & \nabla_{s_1} Q^w(s_1, \mu_{\theta}(s_1)) + \gamma G_1(\mu_{\theta}, s_1), \end{aligned} \quad (10)$$

where  $s_1$  is the next state of  $s$ , which can be generalized to  $k(k \geq 2)$  rollout steps. Let  $s_i$  denote the state visited by the policy at the  $i$ -th step starting from the initial state  $s_0$ ,  $g(s, t, \mu_{\theta}, T) = \prod_{i=1}^{t-1} \nabla_{s_i} T(s_i, \mu_{\theta}(s_i))$ . We choose to

rollout  $k-1$  steps to get rewards, then replace  $\nabla_{s_k} V^{\mu_{\theta}}(s_k)$  by  $\nabla_{s_k} Q^w(s_k, \mu_{\theta}(s_k))$  in Eq. (9), and we get

$$\begin{aligned} L_k(\mu_{\theta}, s, r, T) = & \sum_{t=1}^{k-1} \gamma^{t-1} g(s, t, \mu_{\theta}, T) \nabla_{s_t} r(s_t, \mu_{\theta}(s_t)) \\ & + \gamma^{k-1} g(s, k, \mu_{\theta}, T) \nabla_{s_k} Q^w(s_k, \mu_{\theta}(s_k)). \end{aligned}$$

Replacing  $\nabla_{s'} V^{\mu_{\theta}}(s')$  with  $L_k(\mu_{\theta}, s, r, T)$  in Eq. (8), we get a  $k$ -step estimator of the value gradients:

$$\begin{aligned} G_k(\mu_{\theta}, s) = & \nabla_{\theta} r(s, \mu_{\theta}(s)) + \gamma \nabla_{\theta} T(s, \mu_{\theta}(s)) \\ & L_k(\mu_{\theta}, s, r, T) + \gamma G_k(\mu_{\theta}, s_1). \end{aligned} \quad (11)$$

It is easy to see that  $G_k(\mu_{\theta}, s)$  and  $G_1(\mu_{\theta}, s)$  are the same if we have the true reward and transition functions, which is generally not the case as we need to learn the model in practical environments.

We estimate the recursive value gradient form by the same way as the implementation of the DDPG algorithm. Let  $D_k(\mu_{\theta}, s, T, r)$  denote the value gradient at the sampled state  $s$  with  $k$  rollout steps, on the true transition function  $T$  and reward function  $r$ , which is defined as:

$$\begin{aligned} D_k(\mu_{\theta}, s, T, r) = & \nabla_{\theta} r(s, \mu_{\theta}(s)) + \gamma \nabla_{\theta} T(s, \mu_{\theta}(s)) \\ & L_k(\mu_{\theta}, s, r, T). \end{aligned} \quad (12)$$

To be specific, the recursive value gradient form  $G_k(\mu_{\theta}, s_0)$  can be rewritten as  $\sum_{j=0}^{\infty} \gamma^j D_k(\mu_{\theta}, s_j, T, r)$ . Let  $\rho(s_0, s)$  denote the discounted state distribution from state  $s_0$ , then  $\sum_{j=0}^{\infty} \gamma^j D_k(\mu_{\theta}, s_j, T, r)$  can be rewritten as  $\sum_s \rho(s_0, s) D_k(\mu_{\theta}, s, T, r)$ .

Then, we propose the deterministic value gradients with infinite horizon, where the algorithm is shown in Algorithm 1. As the model is not given, we choose to predict the reward function and the transition function. Given  $n$  samples  $(s_j, a_j, r_j, s_{j+1})$ , for each choice of  $k$ , we employ the empirical state distribution to approximate the discounted state distribution. With the learned transition function  $T'$  and reward function  $r'$ , we use  $\frac{1}{n} \sum_j D_k(\mu_{\theta}, s_j, T', r')$  to update the current policy. We choose to use experience replay to compare with the DDPG algorithm fairly. Different choices of the number of rollout steps trade-off between the variance and the bias. Larger steps means lower variance of the value gradients, and larger bias due to the accumulated model error.

### The Difference between Infinite and Finite Horizon

In this section, we discuss the advantage of our proposed DVG algorithm over finite horizon and validate the effect on a continuous control task. The estimator of deterministic value gradients with finite horizon,  $DVG_F$ , is defined as (Fairbank and Alonso 2012):

$$\begin{aligned} F_k(\mu_{\theta}, s) = & \nabla_{\theta} r(s, \mu_{\theta}(s)) + \gamma \nabla_{\theta} T(s, \mu_{\theta}(s)) \sum_{t=1}^{k-1} \gamma^{t-1} \\ & g(s, t, \mu_{\theta}, T) \nabla_{s_t} r(s_t, \mu_{\theta}(s_t)) + \gamma F_k(\mu_{\theta}, s_1). \end{aligned}$$




**Algorithm 1** The DVG( $k$ ) algorithm

- 1: Initialize the reward network  $r'$ , transition network  $T'$ , critic network  $Q$ , actor network  $\mu_\theta$ , target networks  $Q'$ ,  $\mu'_\theta$  and experience replay buffer  $\mathcal{B}$
- 2: **for** episode = 0, ...,  $N - 1$  **do**
- 3:   **for**  $t = 1, \dots, T$  **do**
- 4:     Select action according to the current policy and exploration noise
- 5:     Execute action  $a_t$ , observe reward  $r_t$  and new state  $s_{t+1}$ , and store transition  $(s_t, a_t, r_t, s_{t+1})$  in  $\mathcal{B}$
- 6:     Sample a random minibatch of  $n$  transitions from  $\mathcal{B}$
- 7:     Update the critic  $Q$  by minimizing the TD error:  $\frac{1}{n} \sum_j (r_j + \gamma Q'(s_{j+1}, \mu_\theta(s_{j+1})) - Q(s_j, a_j))^2$
- 8:     Update the reward network  $r'$  and the transition network  $T'$  on the batch by minimizing the square loss
- 9:     Estimate the value gradients by  $\frac{1}{n} \sum_j D_k(\mu_\theta, s_j, T', r')$  and perform gradient update on the policy
- 10:    Update the target networks by  $\theta^{Q'} = \tau \theta^Q + (1 - \tau) \theta^{Q'}$ ;  $\theta^{\mu'} = \tau \theta^\mu + (1 - \tau) \theta^{\mu'}$
- 11:   **end for**
- 12: **end for**

Note that  $F_k(\mu_\theta, s)$  does not take rewards after the  $k$ -th step into consideration. Therefore, given  $n$  samples  $\{(s_j, a_j, r_j, s_{j+1})\}$ , DVG<sub>F</sub> uses the sample mean of  $D'_k(\mu_\theta, s, T', r')$  to update the policy, where  $D'_k(\mu_\theta, s, T', r')$  is defined as:

$$D'_k(\mu_\theta, s, T', r') = \nabla_\theta r'(s, \mu_\theta(s)) + \gamma \nabla_\theta T'(s, \mu_\theta(s)) \sum_{t=1}^{k-1} \gamma^{t-1} g(s, t, \mu_\theta, T') \nabla_{s'_t} r'(s'_t, \mu_\theta(s'_t)).$$

We then test the two approaches on the environment HumanoidStandup-v2, where we choose the parameter  $k$  to be  $2^4$ . As shown in Figure 2, DVG significantly outperforms DVG<sub>F</sub>, which validates our claim that only considering finite horizon fails to achieve the same performance as that of infinite horizon.

Figure 2: Comparisons of DVG and DVG<sub>F</sub>.

<sup>4</sup>For the choice of  $k$ , we test DVG<sub>F</sub> with steps ranging from 1 to 5, and we choose the parameter with the best performance for fair comparison.




Figure 3: Comparisons of DVG with DDPG.

**Connection and Comparison of DVG and DDPG**

By the proof of the DPG theorem in (Silver et al. 2014), Eq. (8) can be re-written as

$$\nabla_\theta V^{\mu_\theta}(s) = \nabla_\theta \mu_\theta(s) \nabla_a Q^{\mu_\theta}(s, a) + \gamma \nabla_\theta V^{\mu_\theta}(s_1). \quad (13)$$

The DDPG algorithm uses the gradient of the estimator of the  $Q$  function over the action,  $\nabla_a Q^w(s, a)$  to estimate  $\nabla_a Q^{\mu_\theta}(s, a)$ , i.e.,  $G_0(\mu_\theta, s) = \nabla_\theta \mu_\theta(s) \nabla_a Q^w(s, a) + \gamma G_0(\mu_\theta, s_1)$ .

The DDPG algorithm is a model-free algorithm which does not predict the reward and the transition, and can be viewed as the DVG(0) algorithm. We compare the DVG algorithm with different rollout steps  $k$  and DDPG on a continuous control task in MuJoCo, Hopper-v2. From Figure 3, we get that DVG with any choice of the number of rollout steps is more sample efficient than DDPG, which validates the power of model-based techniques. DVG(1) outperforms DDPG and DVG with other number of rollout steps in terms of performance as it trades off well between the bias and the variance of the value gradients. Note that with a larger number of step, DVG(5) is not stable due to the propagated model error.

**The DVPG Algorithm**

As discussed before, the model-based DVG algorithm are more sample efficient than the model-free DDPG algorithm. However, it suffers from the model bias which results in performance loss. In this section, we consider to ensemble these different gradient estimators for better performance.

Motivated by the idea of TD( $\lambda$ ) algorithm (Sutton and Barto 2018), which ensembles the TD( $k$ ) error with a geometric decaying weight  $\lambda$ , we propose a temporal-difference method to ensemble DVG with varying rollout steps and the model-free deterministic policy gradients. We defined the temporal difference deterministic value gradients as  $G_{\lambda,t}(\mu_\theta, s) = (1 - \lambda) \sum_{k=0}^t \lambda^k G_k(\mu_\theta, s)$ , where  $t$  denotes the maximal number of rollout steps by the learned model. For the gradient update rule, we also apply sample based methods: given  $n$  samples  $\{(s_j, a_j, r_j, s_{j+1})\}$ , we use

$$\frac{1}{n} \sum_j ((1 - \lambda) \nabla_\theta \mu_\theta(s_j) \nabla_a Q^w(s_j, a) + (1 - \lambda) \sum_{k=1}^t \lambda^k D_k(\mu_\theta, s_j, T', r')) \quad (14)$$

to update the policy. Based on this ensembled deterministic value-policy gradients, we propose the deterministic value-policy gradient algorithm, shown in Algorithm 2<sup>5</sup>.

---

**Algorithm 2** The DVPG algorithm

---

- 1: Initialize the weight  $\lambda$  and the number of maximal steps  $t$
  - 2: Initialize the reward network  $r'$ , transition network  $T'$ , critic network  $Q$ , actor network  $\mu_\theta$ , target networks  $Q'$ ,  $\mu'_\theta$  and experience replay buffer  $\mathcal{B}$
  - 3: **for** episode = 0, ...,  $N - 1$  **do**
  - 4:   **for**  $t = 1, \dots, T$  **do**
  - 5:     Select action according to the current policy and exploration noise
  - 6:     Execute action  $a_t$ , observe reward  $r_t$  and new state  $s_{t+1}$ , and store transition  $(s_t, a_t, r_t, s_{t+1})$  in  $\mathcal{B}$
  - 7:     Sample a random minibatch of  $n$  transitions from  $\mathcal{B}$
  - 8:     Update the critic  $Q$  by minimizing the TD error:  $\frac{1}{n} \sum_j (r_j + \gamma Q'(s_{j+1}, \mu_\theta(s_{j+1})) - Q(s_j, a_j))^2$
  - 9:     Update the reward network  $r'$  and the transition network  $T'$  on the batch by minimizing the square loss
  - 10:    Estimate the value gradients by Eq. (14), and perform gradient update on the policy
  - 11:    Update the target networks by  $\theta^{Q'} = \tau \theta^Q + (1 - \tau) \theta^{Q'}$ ;  $\theta^{\mu'} = \tau \theta^\mu + (1 - \tau) \theta^{\mu'}$
  - 12:   **end for**
  - 13: **end for**
- 

## Experimental Results

We design a series of experiments to evaluate DVG and DVPG. We investigate the following aspects: (1) What is the effect of the discount factor on DVG? (2) How sensitive is DVPG to the hyper-parameters? (3) How does DVPG compare with state-of-the-art methods?

We evaluate DVPG in a number of continuous control benchmark tasks in OpenAI Gym based on the MuJoCo simulator. We compare DVPG with DDPG, DVG, DDPG with imagination rollouts (DDPG(model)), and SVG with 1 step rollout and experience replay (SVG(1)) in the text. We also compare DVPG with methods using stochastic policies, e.g. ACKTR, TRPO, and DVPG outperforms ACKTR and TRPO significantly. We omit the results due to the limit of the space. We plot the averaged rewards of episodes over 5 different random seeds with the number of real samples, and the shade region represents the 75% confidence interval. We choose the same hyperparameters of the actor and critic network for all algorithms. The prediction models of DVPG, DVG and DDPG(model) are the same.

### The Effect of Discount Factors on DVG

From Eq. (9), we get that  $\nabla_s V^{\mu_\theta}(s)$  is equivalent to the infinite sum of the gradient vectors. To study the effect of

<sup>5</sup>The only difference between the DVG(k) algorithm and the DVPG algorithm is the update rule of the policy.

the discount factor on DVG, we train the algorithm with 2 rollout steps with different values of the discount factor on the environment InvertedPendulum-v2. As shown in Figure 5, 0.95 performs the best in terms of rewards and stability while 0.85 and 0.99 performs comparably, while the performance of 0.8 and 0.6 are inferior to other values. This is because the convergence of the computation of the gradient of the value function over the state may be slow if the discount factor is close to 1 while a smaller value of  $\gamma$  may enable better convergence of  $\nabla_s V^{\mu_\theta}(s)$ . However, the sum of rewards discounted by a too small  $\gamma$  will be too myopic, and fails to perform good. Here, 0.95 trades-off well between the stability and the performance, which is as expected that there exists an optimal intermediate value for the discount factor.




Figure 5: The effect of discount factors.

### Ablation Study of DVPG

We evaluate the effect of the weight of bootstrapping on DVPG with different values from 0.1 to 0.9, where the number of rollout steps is set to be 4. From Figure 6, we get that the performance of the DVPG decreases with the increase of the value  $\lambda$ , where 0.1 performs the best in terms of the sample efficiency and the performance. Thus, we choose the value of the weight to be 0.1 in all experiments.




Figure 6: The weight of bootstrapping.

We evaluate the effect of the number of rollout steps ranging from 1 to 5. Results in Figure 7 show that DVPG with different number of rollout steps all succeed to learn a good policy, with 1 rollout step performing the best. Indeed, the number of rollout steps trade off between the model-error and the variance. There is an optimal value of the number of rollout steps for each environment, which is the only one parameter we tune. To summarize, for the number of look steps,

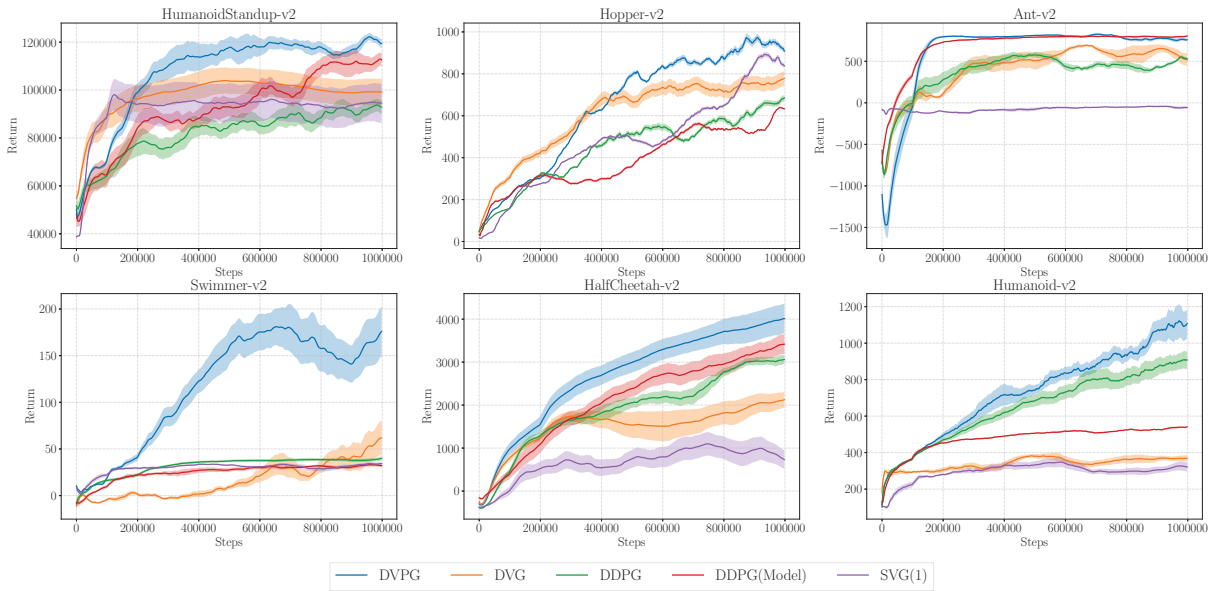


Figure 4: Performance comparisons on environments from the MuJoCo simulator.

1 rollout step works the best on Humanoid-v2, Swimmer-v2 and HalfCheetah-v2, while 2 rollout steps performs the best on HumanoidStandup-v2, Hopper-v2 and Ant-v2. For fair comparisons, we choose the same number of rollout steps for both the DVG and the DVPG algorithm.




Figure 7: The number of rollout steps.

## Performance Comparisons

In this section we compare DVPG with the model-free baseline DDPG, and model-based baselines including DVG, DDPG(model) and SVG(1) on several continuous control tasks on MuJoCo. As shown in Figure 4, there are two classes of comparisons.

Firstly, we compare DVPG with DDPG and DVG to validate the effect of the temporal difference technique to ensemble model-based and model-free deterministic value gradients. The DVG algorithm is the most sample efficient than other algorithms in environments HumanoidStandup-v2, and Hopper-v2. For sample efficiency, DVPG outperforms DDPG as it trades off between the model-based deterministic value gradients and the model-free deterministic policy gradients. In the end of the training, DVPG outperforms other two al-

gorithms significantly, which demonstrates the power of the temporal difference technique. In other four environments, DVPG outperforms other algorithms in terms of both sample efficiency and performance. The performance of DVG and DDPG on Swimmer-v2 and Ant-v2 are comparable, while DVG performs bad in Halfcheetah-v2 and Humanoid-v2 due to the model-error.

Secondly, we compare DVPG with SVG(1) and DDPG with imagination rollouts. Results show that the DVPG algorithm significantly outperforms these two model-based algorithms in terms of sample efficiency and performance, especially in environments where other model-based algorithms do not get better performance than the model-free DDPG algorithm. For the performance of the SVG(1) algorithm, it fails to learn good policies in Ant-v2, which is also reported in (Kurutach et al. 2018).

## Conclusion

Due to high sample complexity of the model-free DDPG algorithm and high bias of the deterministic value gradients with finite horizon, we study the deterministic value gradients with infinite horizon. We prove the existence of the deterministic value gradients for a certain set of discount factors in this infinite setting. Based on this theoretical guarantee, we propose the DVG algorithm with different rollout steps by the model. We then propose a temporal difference method to ensemble deterministic value gradients and deterministic policy gradients, to trade off between the bias due to the model error and the variance of the model-free policy gradients, called the DVPG algorithm. We compare DVPG on several continuous control benchmarks. Results show that DVPG substantially outperforms other baselines in terms of convergence and performance. For future work, it is promising to apply the temporal difference technique presented in

this paper to other model-free algorithms and combine with other model-based techniques.

## Acknowledgments

The work by Ling Pan was supported in part by the National Natural Science Foundation of China Grants 61672316. Pingzhong Tang was supported in part by the National Natural Science Foundation of China Grant 61561146398, and a China Youth 1000-talent program.

## References

- Buckman, J.; Hafner, D.; Tucker, G.; Brevdo, E.; and Lee, H. 2018. Sample-efficient reinforcement learning with stochastic ensemble value expansion. In *Advances in Neural Information Processing Systems*, 8234–8244.
- Cai, Q.; Filos-Ratsikas, A.; Tang, P.; and Zhang, Y. 2018. Reinforcement mechanism design for e-commerce. In *Proceedings of the 2018 World Wide Web Conference*, 1339–1348. International World Wide Web Conferences Steering Committee.
- Chua, K.; Calandra, R.; McAllister, R.; and Levine, S. 2018. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems*, 4754–4765.
- Deisenroth, M., and Rasmussen, C. E. 2011. Pilco: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on machine learning (ICML-11)*, 465–472.
- Fairbank, M., and Alonso, E. 2012. Value-gradient learning. In *Neural Networks (IJCNN), The 2012 International Joint Conference on*, 1–8. IEEE.
- Fairbank, M. 2008. Reinforcement learning by value gradients. *arXiv preprint arXiv:0803.3539*.
- Feinberg, V.; Wan, A.; Stoica, I.; Jordan, M. I.; Gonzalez, J. E.; and Levine, S. 2018. Model-based value estimation for efficient model-free reinforcement learning. *arXiv preprint arXiv:1803.00101*.
- Gu, S.; Lillicrap, T.; Sutskever, I.; and Levine, S. 2016. Continuous deep q-learning with model-based acceleration. In *International Conference on Machine Learning*, 2829–2838.
- Ha, D., and Schmidhuber, J. 2018. Recurrent world models facilitate policy evolution. In *Advances in Neural Information Processing Systems*, 2450–2462.
- Hafner, D.; Lillicrap, T.; Fischer, I.; Villegas, R.; Ha, D.; Lee, H.; and Davidson, J. 2018. Learning latent dynamics for planning from pixels. *arXiv preprint arXiv:1811.04551*.
- Heess, N.; Wayne, G.; Silver, D.; Lillicrap, T.; Erez, T.; and Tassa, Y. 2015. Learning continuous control policies by stochastic value gradients. In *Advances in Neural Information Processing Systems*, 2944–2952.
- Jordan, M. I., and Rumelhart, D. E. 1992. Forward models: Supervised learning with a distal teacher. *Cognitive science* 16(3):307–354.
- Kurutach, T.; Clavera, I.; Duan, Y.; Tamar, A.; and Abbeel, P. 2018. Model-ensemble trust-region policy optimization. *arXiv preprint arXiv:1802.10592*.
- Levine, S., and Koltun, V. 2013. Guided policy search. In *International Conference on Machine Learning*, 1–9.
- Li, W., and Todorov, E. 2004. Iterative linear quadratic regulator design for nonlinear biological movement systems. In *ICINCO (1)*, 222–229.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2015. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*.
- Maaten, L. v. d., and Hinton, G. 2008. Visualizing data using t-sne. *Journal of machine learning research* 9(Nov):2579–2605.
- Moldovan, T. M.; Levine, S.; Jordan, M. I.; and Abbeel, P. 2015. Optimism-driven exploration for nonlinear systems. In *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, 3239–3246. IEEE.
- Montgomery, W., and Levine, S. 2016. Guided policy search as approximate mirror descent. *arXiv preprint arXiv:1607.04614*.
- Nguyen, D. H., and Widrow, B. 1990. Neural networks for self-learning control systems. *IEEE Control systems magazine* 10(3):18–23.
- Oh, J.; Singh, S.; and Lee, H. 2017. Value prediction network. *arXiv preprint arXiv:1707.03497*.
- Schulman, J.; Levine, S.; Abbeel, P.; Jordan, M.; and Moritz, P. 2015. Trust region policy optimization. In *International Conference on Machine Learning*, 1889–1897.
- Silver, D.; Lever, G.; Heess, N.; Degris, T.; Wierstra, D.; and Riedmiller, M. 2014. Deterministic policy gradient algorithms. In *International Conference on Machine Learning (ICML)*.
- Sutton, R. S., and Barto, A. G. 2018. *Reinforcement learning: An introduction*.
- Sutton, R. S.; McAllester, D. A.; Singh, S. P.; and Mansour, Y. 2000. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, 1057–1063.
- Tang, P. 2017. Reinforcement mechanism design. In *IJCAI*, volume 17, 26–30.
- Wahlström, N.; Schön, T. B.; and Deisenroth, M. P. 2015. From pixels to torques: Policy learning with deep dynamical models. *arXiv preprint arXiv:1502.02251*.
- Wang, X., and Dietterich, T. G. 2003. Model-based policy gradient reinforcement learning. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, 776–783.
- Watter, M.; Springenberg, J.; Boedecker, J.; and Riedmiller, M. 2015. Embed to control: A locally linear latent dynamics model for control from raw images. In *Advances in neural information processing systems*, 2746–2754.
- Werbos, P. J. 1990. A menu of designs for reinforcement learning over time. *Neural networks for control* 67–95.