
Trajectory-Based Off-Policy Deep Reinforcement Learning

Andreas Doerr^{1 2 3} Michael Volpp¹ Marc Toussaint³ Sebastian Trimpe²
Christian Daniel¹

Abstract

Policy gradient methods are powerful reinforcement learning algorithms and have been demonstrated to solve many complex tasks. However, these methods are also data-inefficient, afflicted with high variance gradient estimates, and frequently get stuck in local optima. This work addresses these weaknesses by combining recent improvements in the reuse of off-policy data and exploration in parameter space with deterministic behavioral policies. The resulting objective is amenable to standard neural network optimization strategies like stochastic gradient descent or stochastic gradient Hamiltonian Monte Carlo. Incorporation of previous rollouts via importance sampling greatly improves data-efficiency, whilst stochastic optimization schemes facilitate the escape from local optima. We evaluate the proposed approach on a series of continuous control benchmark tasks. The results show that the proposed algorithm is able to successfully and reliably learn solutions using fewer system interactions than standard policy gradient methods.

1. Introduction

Policy search methods are amongst the few successful Reinforcement Learning (RL) (Sutton et al., 2000) methods which are applicable to high-dimensional or continuous control problems, such as the ones typically encountered in robotics (Peters & Schaal, 2008b; Deisenroth et al., 2013). One particular class of policy search methods directly estimates the gradient of the expected return with respect to the parameters of a differentiable policy. These Policy Gradient (PG) algorithms have achieved impressive results on highly complex tasks (Schulman et al., 2015; 2017). However,

standard algorithms are vastly data-inefficient and rely on millions of data points to achieve the aforementioned results. Typical applications are therefore limited to simulated problems where policy rollouts can be cheaply obtained.

Algorithms based on stochastic policy gradients, like REINFORCE (Williams, 1992) and G(PO)MDP (Baxter & Bartlett, 2001), typically estimate the policy gradient based on a batch of trajectories, which are obtained by executing the current policy on the system (i.e. based on on-policy samples). In the next step, all previous experience is discarded and new trajectories are sampled using the updated policy. This scheme holds true also for more recent methods, like PPO (Schulman et al., 2017) or POIS (Metelli et al., 2018), where a surrogate objective is constructed, which can be optimized till convergence. Typically, Importance Sampling (IS) techniques are employed to evaluate a target policy based on rollouts obtained from behavioural policies (i.e. from off-policy samples). Albeit these off-policy evaluation schemes, in these algorithms, no data is shared between iterations. Prominent examples of off-policy offline algorithms typically employ actor-critic architectures (Silver et al., 2014), where the parametric critic model, typically a value function, is updated to summarize all knowledge gathered so far. In contrast, we proposed the model-free Deep Deterministic Off-Policy Gradient method (DD-OPG)¹, which incorporates previously gathered rollout data by sampling from a trajectory replay buffer. This effectively enables backtracking to promising solutions, whilst requiring only minimal assumptions to construct the surrogate model.

Next to the inefficient use of available data, stochasticity in both the policy and the environment causes highly variable gradient estimates and therefore slow convergence. When executing the probabilistic policy on the system, noise is injected into the policy gradient in each time step, leading to a variance, which linearly increases with the length of the horizon (Munos, 2006). Additive Gaussian noise is typically employed as source of exploration. Additionally, PG methods built around the likelihood ratio trick intrinsically require probabilistic policies. Only then, policies can be updated to increase the likelihood of actions, which have been

¹Bosch Center for Artificial Intelligence, Renningen, Germany.
²Max Planck Institute for Intelligent Systems, Stuttgart/Tübingen, Germany. ³Machine Learning and Robotics Lab, University of Stuttgart, Germany. Correspondence to: Andreas Doerr <andreas.doerr@gmx.net>.

¹https://github.com/boschresearch/DD_OPG

advantageous in previous rollouts. Instead of independent noise, temporally-correlated noise (Osband et al., 2016), or exploration directly in parameter space can lead to a larger variety of behaviours (Plappert et al., 2017). Here, the behavioural policy is deterministic, thereby effectively reducing the gradient variance. Methods like DPG (Silver et al., 2014) and DDPG (Lillicrap et al., 2015) learn a parametric value function model to translate changes in policy and therefore actions to changes in expected value. Similarly, our proposed model-free DD-OPG algorithm constructs a non-parametric critic based on importance sampling. This critic, called surrogate model in the following, allows for updating a deterministic policy without the need for explicit parametric value models.

To summarize: We propose an importance sampling based surrogate model of the return distribution, which enables off-policy, offline policy optimization. This surrogate facilitates deterministic policy gradients to reduce gradient variance and enables incorporation of all available data from a replay buffer. Exploration in the policy parameter space is achieved by a prioritized resampling of the surrogates support data, thus favouring promising regions in policy space. Normalized IS, which we demonstrate to act similarly as a baseline in standard PG methods, additionally reduces the variance of the employed estimates. Although no additional, parametric value function baseline (as utilized in TRPO/PPO for variance reduction) is required in our method, fast progress and therefore data-efficient learning is demonstrated on typical continuous control tasks.

The general problem formulation and policy gradient framework is highlighted in Sec. 2, followed by a short presentation of the standard importance sampling estimators to incorporate off-policy data in Sec. 3. The surrogate model, necessary to efficiently incorporate deterministic policy data, as the core of the proposed model-free DD-OPG method is detailed in Sec. 4. In Sec. 5, the main policy optimization scheme is presented and experimentally evaluated in Sec. 6. This work closes with a discussion of connections to related work in Sec. 7 and concludes with an outlook into future work and open topics in Sec. 8.

2. Preliminaries

This section depicts the general episodic RL problem in a discrete-time Markovian environment and summarizes as core building-block of the proposed DD-OPG method, the standard return based policy gradient estimators (Williams, 1992). DD-OPG closely follows this algorithmic structure (cf. Alg. 1), however with extensions to incorporate deterministic, off-policy rollouts as detailed in the following sections. The RL problem is characterized by a discrete-time Markov Decision Process (MDP) $\mathcal{M} = (\mathcal{S}, \mathcal{A}, p, r, \gamma, p_0)$. An agent is interacting with an environment, whose states

$s_t \in \mathcal{S}$ transitions according to the agent’s actions $a_t \in \mathcal{A}$ and the environment’s transition probabilities $p(s_{t+1} | s_t, a_t)$ into a successor state. Starting from a state s_0 drawn from the initial state distribution $p(s_0)$, agent tries to maximize its discounted reward, according to a reward function $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ and discount factor γ , accumulated over a horizon length H . In policy search, the agent acts according to a (stochastic) policy $\pi_\theta = \pi(a_t | s_t; \theta)$, parameterized by θ . The expected accumulated reward is given by

$$J(\theta) = \int p(\tau | \theta) R(\tau) d\tau, \quad (1)$$

where the trajectory $\tau \in \mathcal{T}$ is the sequence of state-action pairs $\tau = (s_0, a_0, \dots, s_H, a_H)$, the (discounted) trajectory return is given by $R(\tau) = \sum_{t=0}^{H-1} \gamma^t r(s_{\tau,t}, a_{\tau,t})$, and due to the Markov property, the trajectory distribution in (1) is given by

$$p(\tau | \theta) = p(s_0) \prod_{t=0}^{H-1} p(s_{t+1} | s_t, a_t) \pi(a_t | s_t; \theta). \quad (2)$$

The dynamics of the system $p(s_{t+1} | s_t, a_t)$ and the initial state distribution $p(s_0)$ are generally unknown to the learning agent.

Model-free policy gradient methods typically directly estimate the expected cost gradient based on the log-derivative trick. The gradient is given by

$$\nabla_\theta J(\theta) = \int p(\tau | \theta) \nabla_\theta \log p(\tau | \theta) R(\tau) d\tau. \quad (3)$$

Given on-policy samples $\tau_i \sim p(\tau | \theta)$, the following Monte Carlo (MC) estimators are obtained for the expected return

$$\hat{J}^{MC}(\theta) = \frac{1}{N} \sum_{i=1}^N R(\tau_i), \quad (4)$$

and the policy gradient

$$\nabla_\theta \hat{J}^{MC}(\theta) = \frac{1}{N} \sum_{i=1}^N \left[\sum_{t=0}^{H-1} \nabla_\theta \log \pi(a_t | s_t; \theta) R(\tau_i) \right]. \quad (5)$$

Since the unknown initial state and dynamics distributions are independent of the policy parameters θ (cf. (2)), the trajectory likelihood gradient $\nabla_\theta \log p(\tau | \theta)$ with respect to the policy parameters can be computed analytically for a given, differentiable policy $\nabla_\theta \log \pi(a_t | s_t; \theta)$.

3. Off-Policy Evaluation

The MC estimators require a substantial amount of on-policy rollouts $\tau_i \sim p(\tau | \theta^*)$ to reduce the gradient estimator’s variance and typically many more rollouts than used in state-of-the-art implementations to closely approximate the true gradient (Ilyas et al., 2018).

For off-policy data, Importance Sampling (IS) can be utilized to incorporate trajectories from a behavioural policy $\pi_{\theta'}$ in order to evaluate a new target policy π_{θ^*} (Zhao et al., 2013; Espeholt et al., 2018; Munos et al., 2016; Metelli et al., 2018). In general, a Monte Carlo estimate of an expectation $\int p(x)f(x)dx$ (such as (1)) can be obtained by sampling from a tractable distribution $x_i \sim q(x)$ and re-weighting the sampled function evaluations $f(x_i)$ based on the likelihood-ratio $p(x_i)/q(x_i)$. The expected return can be rewritten as

$$J(\theta) = \int p(\tau | \theta') \frac{p(\tau | \theta)}{p(\tau | \theta')} R(\tau) d\tau, \quad (6)$$

such that the IS weighted Monte Carlo estimator is given by

$$\hat{J}^{IS}(\theta) = \frac{1}{N} \sum_{i=0}^N \frac{p(\tau_i | \theta)}{p(\tau_i | \theta')} R(\tau_i) \quad (7)$$

$$= \frac{1}{N} \sum_{i=0}^N w(\tau_i, \theta) R(\tau_i), \quad (8)$$

where N trajectories are sampled from a policy $\pi_{\theta'}$ to infer the expected cost of policy π_{θ} . Although system dynamics and initial state distribution in (2) are unknown, the likelihood-ratio, i.e. the importance weights, can be computed since the unknown parts cancel out, such that

$$w(\tau, \theta) = \frac{p(\tau | \theta)}{p(\tau | \theta')} = \frac{\prod_{t=0}^H \pi(a_t | s_t; \theta)}{\prod_{t=0}^H \pi(a_t | s_t; \theta')}. \quad (9)$$

During learning, trajectories are collected from multiple different policies $\mathcal{D} = \{(\tau_i, \theta_i)\}_{i=1}^N$. To incorporate all data, the importance sampling distribution can be replaced by an empirical mixture distribution $q(\tau | \theta_1, \dots, \theta_N) = 1/N \sum_i p(\tau | \theta_i)$ such that the available trajectories are i.i.d. draws from the empirical mixture distribution $\tau_i \sim q(\tau | \theta_1, \dots, \theta_N)$ (Jie & Abbeel, 2010). The resulting importance weights are given by

$$w(\tau, \theta) = \frac{\prod_{t=0}^H \pi(a_t | s_t; \theta)}{\frac{1}{N} \sum_j \prod_{t=0}^H \pi(a_t | s_t; \theta_j)}. \quad (10)$$

Computing the importance weights in (10), however, scales quadratically with the number of available trajectories due to the summation over the likelihoods of all trajectories given all available policies. Scaling this estimator to today's deep neural network policies with a large number of required rollouts is, thus, a major challenge. Instead of computing the surrogate based on all data, as in (Jie & Abbeel, 2010), which is only feasible for several hundred rollouts, the proposed DD-OPG method employs a trajectory replay buffer and a probabilistic selection scheme to recompute a stochastic approximation of the full surrogate model. This idea is related to prioritized experience replay (Schaul et al., 2015) but for full trajectories. It enables scaling to much larger

datasets and at the same time helps to avoid local minima by stochastically optimizing the objective.

Another technique typically employed for IS is weight normalization (Metelli et al., 2018). The *weighted importance sampling* estimator obtains a lower variance estimate at the cost of adding bias. It has been employed in (Peshkin & Shelton, 2002) and is both theoretically and empirically better-behaved (Meuleau et al., 2000; Precup et al., 2000; Shelton, 2001) compared to the pure IS estimator. The weighted importance sampling estimator is given by

$$\hat{J}^{WIS}(\theta) = \frac{1}{Z} \sum_{i=0}^N w(\tau_i, \theta) R(\tau_i), \quad (11)$$

where importance weights $w(\tau_i, \theta)$ might be computed according to (9) or (10) and a normalizing constant $Z = \sum_{i=0}^N w(\tau_i, \theta)$ instead of the standard normalization $Z = N$, previously used in (8).

From the policy gradient perspective, by normalizing the importance weights, we obtain a gradient estimator, which includes a parameter dependent baseline.

Proposition 1 *The policy gradient estimator obtained from the self-normalized importance sampling expected cost estimator \hat{J}^{WIS} is given by*

$$\nabla_{\theta} \hat{J}^{WIS}(\theta) = \frac{1}{Z} \sum_{i=1}^N \left[w(\tau_i, \theta) \sum_{t=0}^H \left[\nabla_{\theta} \log \pi(a_t^{(i)} | s_t^{(i)}; \theta) \right] \left[R(\tau_i) - \hat{J}^{WIS}(\theta) \right] \right]. \quad (12)$$

A proof of this proposition is shown in Appendix A. This estimator is closely related to standard PG estimators with an added baseline term for variance reduction.

In standard, REINFORCE like, PG methods, two of the most common variance reduction techniques (Greensmith et al., 2004) are: i) incorporation of the reward-to-go for each policy action update instead of the entire Monte Carlo path return; and ii) subtraction of a state dependent baseline term, such as to obtain an estimate of the advantage of the previously taken action. The intuition behind method i) is to reward actions only for rewards obtained after the action took effect, but not for those obtained earlier on. However, to compute the importance weights not for the full trajectory distribution but for each state-action pair individually, the computation of a matrix of size $\mathcal{O}(N^2 H^2)$ would be required. Therefore, the model-free, importance sampling based approaches are typically limited to the path return based estimators. Model-based methods (i.e. a parametric models for the value function) are employed in the cost-to-go estimators. Variance reduction method ii) is automatically obtained by the normalized estimator as shown

in proposition 1, however, in contrast to the bias free value function control variates, at the cost of adding bias. Additionally, optimal baselines to further decrease the variance of the gradient estimator have been derived in (Jie & Abbeel, 2010) and could be incorporated into DD-OPG.

4. Deterministic Policy Gradients

The policy gradient estimators in (5) and (12) rely on a policy distribution $\pi(a_t|s_t; \theta)$ in order to obtain a gradient signal on how to update the policy parameters to increase the likelihood of successful actions. In this situation the, typically Gaussian, additive policy noise acts in two ways, causing exploration and serving as the basis for the estimation of the objective function.

Exploration is being driven directly through noise in the action space, i.e., the policy covariance. While driving exploration through noisy actions will converge in the limit, the resulting explorative behaviour exhibits no temporal correlations, which can make it inefficient.

Estimation of the objective function is typically achieved by reweighting the action distribution according to the policy’s likelihood. Standard policies are given as $\pi(a_t|s_t; \theta) = \mathcal{N}(a_t|\mu_\theta(s_t), \Sigma_\theta)$, where μ_θ is represented by some function approximator parameterized by θ , e.g. a neural network. The additive Gaussian noise covariance is typically a diagonal matrix, parameterized by θ as well. The proposed deterministic policy gradient method strives to separate the exploration and estimation part.

Parameter Space Exploration By utilizing deterministic rollout policies, the only noise introduced into the gradient estimate originates from the stochasticity of the environment and we have to perform exploration in parameter space instead of action space exploration. However, as stated above, parameter based exploration may in many cases be more efficient than exploration in action space, since parameter based exploration will lead to temporally correlated actions which can explore the state space faster. Typically, however, this effect is negated for neural network policies since the parameter space that has to be explored is prohibitively large. Thus, to navigate large parameter spaces efficiently, some approximate evaluation of the cost function (1) is needed.

Trajectory based objective estimate Whilst evaluation of the Monte Carlo based expected cost estimate is possible also for deterministic policies, the off-policy evaluation is no longer feasible since the likelihood ratio $p(\tau|\theta)/p(\tau|\theta')$ (cf. (9)) becomes zero for two distinct dirac policy action distributions if $\mu_\theta(s) \neq \mu_{\theta'}(s)$.

However, we can still compare trajectories under a stochastic evaluation distribution, similar to a kernel function where the standard deviation of the evaluation function relates to a kernel lengthscale in action space.

Thus, we introduce the evaluation policy

$$\tilde{p}(a_t|s_t; \theta) = \mathcal{N}(a_t|\mu_\theta(s_t), \Sigma), \quad (13)$$

where $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_{D_u})$ is a diagonal covariance matrix as typically employed in deep RL methods with Gaussian action noise. The deterministic policy is given by $p(a_t|s_t; \theta) = \delta(a = \mu_\theta(s_t))$, where δ is the dirac delta. From the general IS expectation in (6) and our evaluation policy in (13), the surrogate model follows as

$$\hat{J}^{\text{surr}}(\theta) = \frac{1}{Z} \sum_{i=1}^N \tilde{w}(\tau_i, \theta) R(\tau_i), \quad (14)$$

with surrogate weights

$$\tilde{w}(\tau_i, \theta) = \frac{\prod_{t=0}^H \mathcal{N}(a_t^{(i)}|\mu_\theta(s_t^{(i)}), \Sigma)}{\frac{1}{N} \sum_{j=0}^N \prod_{t=0}^H \mathcal{N}(a_t^{(i)}|\mu_{\theta_j}(s_t^{(i)}), \Sigma)}, \quad (15)$$

where, depending on the choice of normalization constant Z , we obtain the analogue to the standard IS estimator ($Z = N$) or the analog to the weighted IS estimator ($Z = \sum_{i=1}^N \tilde{w}(\tau_i, \theta)$). Reintroducing the fixed Gaussian noise as an implicit loss to obtain gradients for the evaluation of deterministic policies is clearly a model assumption in the proposed method but can be justified from several perspectives.

The hyper-parameter Σ allows for control over the amount of information shared between neighbouring policies. Similar to the cap of importance weights in PPO (Schulman et al., 2017), this parameter allows to control bias and variance of the surrogate model. Analyzing the introduced bias and relation to the PPO weight cap is however ongoing research. In the limit of $\Sigma \rightarrow \mathbf{0}$, the proposed surrogate (14) approaches the MC estimator (4). Only in case of two different policy parameterizations $\theta_j \neq \theta_i$, but equivalent actions $\mu_{\theta_j}(s) = \mu_{\theta_i}(s)$ for the sampled states s , the surrogate model would output an average whereas the MC estimator would not mix up the obtained returns. For $\Sigma = \Sigma_\theta$, the surrogate model recovers the true IS estimate, given that all trajectories are generated using the same additive Gaussian noise. Finally, for $\Sigma \rightarrow \text{inf}$, the estimate is simply the average over all available path returns.

Modelling the expected return distribution by choosing a lengthscale in action space can furthermore be motivated from a second perspective. Typical expected return distributions oftentimes comprise sharp transitions between stable and unstable regions, where policy parameters change only slightly but reward changes drastically. One global lengthscale is therefore typically not well suited to directly model the expected return. This is a standard problem in Bayesian Optimization for reinforcement learning, where typical smooth kernel functions (e.g. squared exponential kernel) with globally fixed lengthscales are unable to model

both stable and unstable regimes at the same time. However, in the proposed model, a lengthscale in action space is translated via the sampled state distribution and policy function μ into implicit assumptions in the actual policy parameter space. Doing so, instead of operating on arbitrary euclidean distances in policy parameter space, a more meaningful distance in trajectory and action space is available. Typically, for a given system, distance of trajectories and between actions is more graspable, compared to arbitrary deep neural network policy parameters.

The expected return estimator (14) falls back to zero for policy evaluation far away from training data. To estimate the variance of the importance sampling estimator itself, typically, the Effective Sample Size (ESS) is evaluated. Based on the variance of the importance weights, it analyses the effective number of available data points at a specific policy evaluation position. In (Metelli et al., 2018), a lower bound on the expected return has been proposed such that with probability $1 - \delta$ it holds that

$$\begin{aligned} \mathbb{E}_{\tau \sim p(\tau|\theta)}[R(\tau)] &\geq \frac{1}{N} \sum_{i=1}^N \tilde{w}(\tau_i, \theta) R(\tau_i) \\ &\quad - \|R\|_\infty \sqrt{\frac{(1-\delta)d_2(p(\tau|\theta)\|p(\tau|\theta'))}{\delta N}}, \end{aligned} \quad (16)$$

where d_2 is the exponentiated 2-Rényi divergence. Due to the identity $\text{ESS}(P||Q) = N/d_2(P||Q)$, this lower bound can be estimated in a sample-based way by employing the ESS estimator

$$\text{ESS} = \frac{1}{\sum_{i=1}^N \tilde{w}(\tau_i, \theta)^2}, \quad (17)$$

such as to obtain the lower bound estimate

$$\begin{aligned} \mathbb{E}_{\tau \sim p(\tau|\theta)}[R(\tau)] &\geq \frac{1}{N} \sum_{i=1}^N \tilde{w}(\tau_i, \theta) R(\tau_i) \\ &\quad - \|R\|_\infty \sqrt{\frac{1-\delta}{\delta} \text{ESS}(\theta)^{-1}}. \end{aligned} \quad (18)$$

Refer to theorem 4.1 in (Metelli et al., 2018) for details and proof regarding the lower bound in (16). The confidence parameter δ determines, similar to the KL-divergence in TRPO (Schulman et al., 2015), how far the policy optimization can step away from known regions. In DD-OPG, this uncertainty estimate is employed as penalty

$$\text{penalty}(\theta) = -\|R\|_\infty \gamma \sqrt{\text{ESS}(\theta)^{-1}}, \quad (20)$$

with penalty factor γ as an hyper-parameter to control exploration, i.e. following the objective estimate vs. risk awareness, i.e. staying within a trust region.

5. Model-Free Off-Policy Optimization

The surrogate model of the return distribution, as derived in Sec. 4, can now be directly incorporated for policy optimization. In related work, parametric search distributions (e.g. Gaussian) are employed as policy search distribution or hyperpolicy (Zhao et al., 2013; Plappert et al., 2017; Metelli et al., 2018). However, in high-dimensional spaces, as typically obtained with deep network policy representations, updating the full search distribution is challenging and common approaches usually revert to heuristics to control a simplified, e.g. diagonal or block-wise search distribution's covariance matrix.

Instead, the proposed model-free DD-OPG method fully optimizes a stochastic version of the surrogate objective to foster exploration and overcome local minima. At the same time, the stochastic evaluation mitigates the unfavourable complexity of computing the full importance sampling estimate based on all available data. Due to the empirical mixture distribution in (10), computing the likelihood of all observed trajectories under all policies is quadratic in the number of observed paths. Instead, the proposed method employs a selection criterion to construct a stochastic surrogate model based on a subset of rollouts in each policy optimization step. In particular, a predefined number of N_{max} rollout indices is drawn from the softmax distribution over the discrete set of available trajectory indices \mathcal{I} . The softmax is computed based on the normalized, empirical returns \tilde{R} and a temperature factor λ .

$$p(\mathcal{I}|\tau_1, \dots, \tau_N) = \frac{\exp(\tilde{R}(\tau_{\mathcal{I}})/\lambda)}{\sum_{j=1}^N \exp(\tilde{R}(\tau_j)/\lambda)}. \quad (21)$$

The temperature λ is used to trade off exploration against exploitation in the selection of reference trajectories. This scheme is closely related to prioritized experience replay (Schaul et al., 2015). A study of the effect of temperature selection on the learning progress is shown in Sec. 6.3.

The full DD-OPG algorithm is detailed in Alg. 1. The main objective is to incorporate all available deterministic policy rollouts, not only the ones from the current iteration, into the surrogate model by means of the softmax replay selection. The lower bound expected return can then be fully optimized using standard optimization techniques. In practice Adam (Kingma & Ba, 2014) is employed, but other techniques, e.g. based on the natural policy gradient (Peters & Schaal, 2008a) could be incorporated as well.

6. Experimental Evaluation

The experimental evaluation of the proposed DD-OPG method is threefold. In Sec. 6.1, the resulting surrogate return model is visualized, highlighting different modeling options. A benchmark against state-of-the-art PG methods

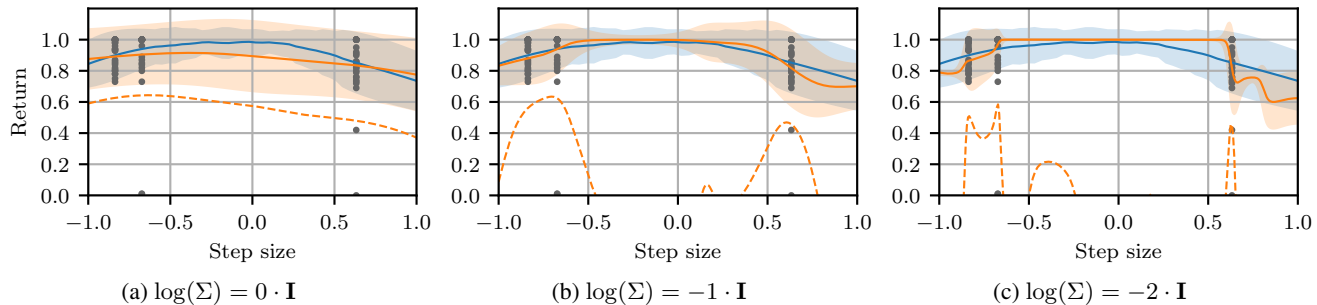


Figure 1. Visualization of the surrogate return model. A cross-section along a random direction in parameter space is shown for parameters close to optimum. The ground truth mean and std (blue) of the return distribution is shown together with the mean and std estimate (orange) from the weighted importance sampling surrogate model. The lower confidence bound ($\delta = 0.2$, dashed orange line) is shown together with the model’s input data (grey dots). Notice how more or less information is shared between points where data is available depending on the chosen lengthscale parameter Σ .

Algorithm 1 Model-free DD-OPG

Input: Initial policy parameters θ_0
 Empty trajectory replay buffer $\mathcal{D}_0 = \{\}$
repeat
 Sample trajectory: $\tau_i \sim p(\tau | \theta_i)$
 Update trajectory buffer: $\mathcal{D}_{i+1} = \mathcal{D}_i \cup (\tau_i, R_i)$
 Memory selection: $i_1, \dots, i_{N_{max}} \stackrel{iid}{\sim} p(\mathcal{I} | \tau_1, \dots, \tau_i)$
 Surrogate model: $\tilde{J}(\theta)$, $\text{penalty}(\theta)$
 Lower bound optimization:
 $\theta_{i+1} = \underset{\theta}{\operatorname{argmax}} \tilde{J}(\theta) - \text{penalty}(\theta)$
until converged or maximum iterations

is shown in Sec. 6.2 to highlight fast and data-efficient learning. Finally, important parts of the proposed algorithms and their effects on the final learning performance are highlighted in an ablation study in Sec. 6.3.

6.1. Surrogate Model

As discussed in Sec. 4, the proposed surrogate model can smoothly interpolate between the Monte Carlo estimate, the importance sampling estimate, and an average of all available returns. In Fig. 1, the available surrogate model predictions are visualized for multiple settings of the model hyper-parameter Σ . In particular, the estimate for expected return (solid orange line), return variance (shaded orange visualizes one standard deviation), and the lower bound of the expected return (dashed orange line) are visualized for policy evaluations along a random direction around the optimal policy θ^* for the cartpole environment (experimental details can be found in Appendix B). Trajectory data, which is available to the estimator is highlighted by grey dots. The groundtruth return distribution (mean \pm one std. in blue) is computed using the standard MC estimator, based on independent policy rollouts, which are *not* part of the surrogate model.

Stepping from long lengthscales (cf. Fig. 1a) to shorter lengthscales (cf. Fig. 1c), the surrogate model predictions become more local. Most visibly in the lowerbound estimate, the ESS drops significantly when moving away from data points and small model lengthscales, resulting in much higher uncertainty.

6.2. Policy Gradient Benchmark

The proposed DD-OPG method is evaluated in terms of data-efficiency and learning progress in comparison to state-of-the-art policy gradient methods based on Monte Carlo return estimates. In contrast, methods such as DDPG (Lillicrap et al., 2015) employ TD learning for their value function model and are not part of this evaluation. The benchmark compares DD-OPG to the standard REINFORCE (Williams, 1992) baseline and both TRPO (Schulman et al., 2015) and PPO (Schulman et al., 2017). All competitor algorithms employ, as it is common practice, the reward-to-go formulation and a linear feature-based baseline for variance reduction. For all methods, hyper-parameters are selected to achieve maximal accumulated average return, i.e. fast and stable policy optimization. Details about the individual methods’ configuration and the employed environments can be found in Appendix B.

The resulting learning performances are visualized in Fig. 2 for the cartpole, mountaincar and swimmer environment (left to right) (Duan et al., 2016). For REINFORCE (blue), TRPO (yellow), PPO (green), and DD-OPG (red), the mean average return (solid line) and its confidence intervals (one standard deviation as shaded area) are depicted, as obtained from 10 independent runs out of 10 random seeds for each environment and method. To compare the learning speed and data-efficiency between the batch-wise learning competitors and the rollout-based DD-OPG, the results are visualized as a function of collected environment interactions (scaled by 10^5) in Fig. 2.




Figure 2. Policy gradient methods benchmark. The proposed method DD-OPG (red) is compared to standard REINFORCE (blue), TRPO (orange) and PPO (green) on three continuous control benchmark problems. Mean and standard deviation of the average return (obtained from 10 independent random seeds) are plotted as a function of the system interaction steps (scaled by 10^5). Significant faster learning speed in the beginning is observed for the model-free off-policy method in comparison to the on-policy PG methods.

With DD-OPG, rapid learning progress is achieved already and the final performance of the competitive, state-of-the-art policy gradient methods is matched. In the hyper-parameter tuning phase, experiments with TRPO and PPO have been conducted based on smaller batchsizes, but due to the lack of data-efficient incorporation of off-policy data, no faster *and* stable learning progress could be achieved for these methods, compared to the one visualized in Fig. 2. Notice the large variance of the DD-OPG learning progress in the swimmer environment. Albeit the superior learning performance of DD-OPG on the swimmer environment, some of the runs got stuck in local minima, resulting in the large variance estimate. This trade-off between exploration and exploitation is partially achieved by the stochastic memory selection. A mix of prioritized trajectory replay and current trajectories is mandatory to prevent greedy exploitation of previously seen, local minima and to facilitate exploration. Our experiments show that it is mandatory to incorporate previously seen rollout data, as it is done in DD-OPG, to enable rapid progress already in the early stages of training.

6.3. Ablation Study

In the final DD-OPG algorithm, multiple aspects come together: i) the deterministic surrogate model, ii) the memory selection strategy, and iii) the optimization scheme. In this ablation study, we separate the individual components to analyse their effect on the final learning performance. Experiments are conducted on the cartpole environment and results are averaged over three random seeds.

In the first experiment, DD-OPG is reconstructed starting from the REINFORCE baseline. A visualization is shown in Fig. 3. In REINFORCE (red dotted line), only one policy gradient step is taken based on the current on-policy data. This is comparable to DD-OPG with almost no memory ($N_{max} = 5$) and only one step gradient update (visualized as blue dotted line). Learning performance is already increased by adding more memory paths (green: $N_{max} = 20$, yellow: $N_{max} = 50$). More significantly, the full opti-




Figure 3. Ablation study of DD-OPG. The full DD-OPG model is constructed from the REINFORCE baseline by iteratively adding i) deterministic off-policy data incorporation and ii) full optimization of the surrogate model. Visualized is the mean learning progress from 3 random seeds on the cartpole environment. REINFORCE (dashed red line) is shown with DD-OPG optimizing only for one gradient step (dotted lines) and fully optimizing the surrogate model (solid lines). For DD-OPG, three levels of history are shown (blue: $N_{max} = 5$, green: $N_{max} = 20$, yellow: $N_{max} = 50$).

mization of the surrogate model (solid lines) achieves much faster learning progress.

In Fig. 4, the effect of the surrogate model’s lengthscale parameter Σ is evaluated. Four different lengthscales $\log \Sigma$ are evaluated (red: 1.0, green: 2.0, yellow: 3.0, blue: 4.0). In this experiment, longer lengthscales clearly improve learning speed despite the introduced model bias.

The effects of the softmax temperature λ on the proposed prioritized trajectory replay and the learning progress are depicted in Fig. 5. Explorative behaviour is favoured for higher temperatures (red), whereas for low temperatures (blue), previous trajectories are selected more greedily. In this example, an intermediate temperature achieves the best trade-off exploration-exploitation trade-off.




Figure 4. Effect of the surrogate hyper-parameter Σ on the learning progress. Learning speed increases from short lengthscales $\log \Sigma = 1.0$ (red) to $\log \Sigma = 2.0$ (green), $\log \Sigma = 3.0$ (yellow), and $\log \Sigma = 4.0$ (blue). Visualized are DD-OPG mean learning curves from three random seeds as a function of the number of interaction steps with the cartpole environment (scaled by 10^3).

7. Connections to Related Work

Policy search methods (Peters & Schaal, 2008b; Deisenroth et al., 2013) and policy gradient methods (Williams, 1992; Baxter & Bartlett, 2001) are well studied in the RL community and many connections to DD-OPG exist.

Importance sampling has been employed to either reweight full trajectory distributions (Shelton, 2001; Jie & Abbeel, 2010; Zhao et al., 2013; Metelli et al., 2018) or to reweight individual state-action pairs (Munos et al., 2016; Espeholt et al., 2018). Except for (Jie & Abbeel, 2010), no global IS estimator is derived, but estimates are only based on the current iteration’s data. In contrast, DD-OPG introduces global surrogate model based on all available deterministic policy rollouts and computes local, stochastic approximations using prioritized replay. Instead of DD-OPG’s action space lengthscale, alternative approaches consider truncation of the importance weights (Wawrzynski & Pacut, 2007; Schulman et al., 2017; Espeholt et al., 2018). So far, the connection between both approaches has not yet been subject of greater analysis.

Concepts for policy updates range from standard gradient ascent (Williams, 1992), to trust region methods (Schulman et al., 2015) to lower bounds, which can be fully optimized till convergence (Schulman et al., 2017; Metelli et al., 2018). The proposed DD-OPG optimizes a stochastic version based on the lower bound, derived in (Metelli et al., 2018).

Deterministic policies as means of variance reduction have been previously discussed for example in (Sehnke et al., 2008; Plappert et al., 2017). Instead of action noise for exploration, exploration is achieved by stochasticity in parameter space. The DD-OPG method relies on deterministic

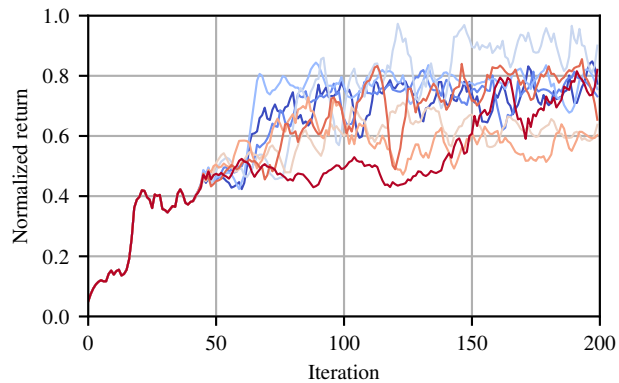


Figure 5. Learning progress for multiple temperature settings for softmax trajectory selection. From lowest temperature ($\lambda = 0.01$, blue) to highest temperature ($\lambda = 2.0$, red). Both too high and too low temperatures lead to suboptimal behaviour, either by too much exploration or too greedy behaviour.

policies for variance reduction, but introduces exploration by means of stochastic gradients from the prioritized replay model.

8. Discussion

This work presents a new surrogate model of the RL return distribution inspired by importance sampling. It can incorporate off-policy data and *deterministic* rollouts to reduce estimator variance. Despite the promising results and the data-efficient learning progress, several interesting topics remain for future work.

The proposed surrogate model is motivated by its close connections to the importance sampling estimator, the interpretability of the model assumption in action space and its desirable behaviour in the model limits. A detailed analysis of the resulting model assumptions in policy space, implied by the model assumptions in action space and an analysis of the resulting bias remains an open question.

The proposed optimization scheme empirically achieved good performance in our benchmark experiments, outperforming state-of-the-art methods, although no additional parametric value function baseline (as in TRPO/PPO) is employed. However, extensions to other strategies for exploration vs. exploitation, for example acquisition functions like Expected Improvement or Probability of Improvement from Bayesian Optimization (Snoek et al., 2012), are to be explored and directly carry over to the proposed surrogate return model.

Finally, memory selection is required to scale the non-parametric model structure to typical deep RL applications. The proposed prioritized trajectory replay is only one possible option to address this challenge.

References

- Baxter, J. and Bartlett, P. L. Infinite-horizon policy-gradient estimation. *Journal of Artificial Intelligence Research*, 15:319–350, 2001.
- Deisenroth, M. P., Neumann, G., Peters, J., et al. A survey on policy search for robotics. *Foundations and Trends® in Robotics*, 2(1–2):1–142, 2013.
- Duan, Y., Chen, X., Houthoofd, R., Schulman, J., and Abbeel, P. Benchmarking deep reinforcement learning for continuous control. In *International Conference on Machine Learning (ICML)*, pp. 1329–1338, 2016.
- Espeholt, L., Soyer, H., Munos, R., Simonyan, K., Mnih, V., Ward, T., Doron, Y., Firoiu, V., Harley, T., Dunning, I., et al. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. In *International Conference on Machine Learning (ICML)*, pp. 1406–1415, 2018.
- Greensmith, E., Bartlett, P. L., and Baxter, J. Variance reduction techniques for gradient estimates in reinforcement learning. *Journal of Machine Learning Research*, 5(Nov):1471–1530, 2004.
- Ilyas, A., Engstrom, L., Santurkar, S., Tsipras, D., Janoos, F., Rudolph, L., and Madry, A. Are deep policy gradient algorithms truly policy gradient algorithms? *arXiv preprint arXiv:1811.02553*, 2018.
- Jie, T. and Abbeel, P. On a connection between importance sampling and the likelihood ratio policy gradient. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 1000–1008, 2010.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- Metelli, A. M., Papini, M., Faccio, F., and Restelli, M. Policy optimization via importance sampling. In *Advances in neural information processing systems (NIPS)*, pp. 5442–5454, 2018.
- Meuleau, N., Peshkin, L., Kaelbling, L. P., and Kim, K.-E. Off-policy policy search. *MIT Artificial Intelligence Laboratory*, 2000.
- Munos, R. Policy gradient in continuous time. *Journal of Machine Learning Research*, 7(May):771–791, 2006.
- Munos, R., Stepleton, T., Harutyunyan, A., and Bellemare, M. Safe and efficient off-policy reinforcement learning. In *Advances in neural information processing systems (NIPS)*, pp. 1054–1062, 2016.
- Osband, I., Blundell, C., Pritzel, A., and Van Roy, B. Deep exploration via bootstrapped DQN. In *Advances in neural information processing systems (NIPS)*, pp. 4026–4034, 2016.
- Peshkin, L. and Shelton, C. R. Learning from scarce experience. In *International Conference on Machine Learning (ICML)*, pp. 498–505. Morgan Kaufmann Publishers Inc., 2002.
- Peters, J. and Schaal, S. Natural actor-critic. *Neurocomputing*, 71(7-9):1180–1190, 2008a.
- Peters, J. and Schaal, S. Reinforcement learning of motor skills with policy gradients. *Neural networks*, 21(4):682–697, 2008b.
- Plappert, M., Houthoofd, R., Dhariwal, P., Sidor, S., Chen, R. Y., Chen, X., Asfour, T., Abbeel, P., and Andrychowicz, M. Parameter space noise for exploration. *arXiv preprint arXiv:1706.01905*, 2017.
- Precup, D., Sutton, R. S., and Singh, S. P. Eligibility traces for off-policy policy evaluation. In *International Conference on Machine Learning (ICML)*, pp. 759–766. Cite-seer, 2000.
- Schaul, T., Quan, J., Antonoglou, I., and Silver, D. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*, 2015.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. Trust region policy optimization. In *International Conference on Machine Learning (ICML)*, pp. 1889–1897, 2015.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Sehnke, F., Osendorfer, C., Rückstieß, T., Graves, A., Peters, J., and Schmidhuber, J. Policy gradients with parameter-based exploration for control. In *International Conference on Artificial Neural Networks*, pp. 387–396. Springer, 2008.
- Shelton, C. R. Policy improvement for POMDPs using normalized importance sampling. In *Proceedings of the Seventeenth conference on Uncertainty in artificial intelligence (UAI)*, pp. 496–503. Morgan Kaufmann Publishers Inc., 2001.

- Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D., and Riedmiller, M. Deterministic policy gradient algorithms. In *ICML*, 2014.
- Snoek, J., Larochelle, H., and Adams, R. P. Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*, pp. 2951–2959, 2012.
- Sutton, R. S., McAllester, D. A., Singh, S. P., and Mansour, Y. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems (NIPS)*, pp. 1057–1063, 2000.
- Wawrzynski, P. and Pacut, A. Truncated importance sampling for reinforcement learning with experience replay. *Proc. CSIT Int. Multiconf.*, pp. 305–315, 2007.
- Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. In *Reinforcement Learning*, pp. 5–32. Springer, 1992.
- Zhao, T., Hachiya, H., Tangkaratt, V., Morimoto, J., and Sugiyama, M. Efficient sample reuse in policy gradients with parameter-based exploration. *Neural computation*, 25(6):1512–1547, 2013.