
Policy Optimization via Importance Sampling

Alberto Maria Metelli

Politecnico di Milano, Milan, Italy
albertomaria.metelli@polimi.it

Matteo Papini

Politecnico di Milano, Milan, Italy
matteo.papini@polimi.it

Francesco Faccio

Politecnico di Milano, Milan, Italy
IDSIA, USI-SUPSI, Lugano, Switzerland
francesco.faccio@mail.polimi.it

Marcello Restelli

Politecnico di Milano, Milan, Italy
marcello.restelli@polimi.it

Abstract

Policy optimization is an effective reinforcement learning approach to solve continuous control tasks. Recent achievements have shown that alternating online and offline optimization is a successful choice for efficient trajectory reuse. However, deciding when to stop optimizing and collect new trajectories is non-trivial, as it requires to account for the variance of the objective function estimate. In this paper, we propose a novel, model-free, policy search algorithm, POIS, applicable in both action-based and parameter-based settings. We first derive a high-confidence bound for importance sampling estimation; then we define a surrogate objective function, which is optimized offline whenever a new batch of trajectories is collected. Finally, the algorithm is tested on a selection of continuous control tasks, with both linear and deep policies, and compared with state-of-the-art policy optimization methods.

1 Introduction

In recent years, policy search methods [10] have proved to be valuable Reinforcement Learning (RL) [50] approaches thanks to their successful achievements in continuous control tasks [e.g., 23, 42, 44, 43], robotic locomotion [e.g., 53, 20] and partially observable environments [e.g., 28]. These algorithms can be roughly classified into two categories: *action-based* methods [51, 34] and *parameter-based* methods [45]. The former, usually known as policy gradient (PG) methods, perform a search in a parametric policy space by following the gradient of the utility function estimated by means of a batch of trajectories collected from the environment [50]. In contrast, in parameter-based methods, the search is carried out directly in the space of parameters by exploiting global optimizers [e.g., 41, 16, 48, 52] or following a proper gradient direction like in Policy Gradients with Parameter-based Exploration (PGPE) [45, 63, 46]. A major question in policy search methods is: how should we use a batch of trajectories in order to exploit its information in the most efficient way? On one hand, *on-policy* methods leverage on the batch to perform a single gradient step, after which new trajectories are collected with the updated policy. Online PG methods are likely the most widespread policy search approaches: starting from the traditional algorithms based on stochastic policy gradient [51], like REINFORCE [64] and G(PO)MDP [4], moving toward more modern methods, such as Trust Region Policy Optimization (TRPO) [42]. These methods, however, rarely exploit the available trajectories in an efficient way, since each batch is thrown away after just one gradient update. On the other hand, *off-policy* methods maintain a behavioral policy, used to explore the environment and to collect samples, and a target policy which is optimized. The concept of off-policy learning is rooted in value-based RL [62, 30, 27] and it was first adapted to PG in [9], using an actor-critic architecture. The approach has been extended to Deterministic Policy Gradient (DPG) [47], which allows optimizing deterministic policies while keeping a stochastic policy for exploration.

More recently, an efficient version of DPG coupled with a deep neural network to represent the policy has been proposed, named Deep Deterministic Policy Gradient (DDPG) [23]. In the parameter-based framework, even though the original formulation [45] introduces an online algorithm, an extension has been proposed to efficiently reuse the trajectories in an offline scenario [67]. Furthermore, PGPE-like approaches allow overcoming several limitations of classical PG, like the need for a stochastic policy and the high variance of the gradient estimates.¹

While on-policy algorithms are, by nature, *online*, as they need to be fed with fresh samples whenever the policy is updated, off-policy methods can take advantage of mixing online and *offline* optimization. This can be done by alternately sampling trajectories and performing optimization epochs with the collected data. A prime example of this alternating procedure is Proximal Policy Optimization (PPO) [44], that has displayed remarkable performance on continuous control tasks. Off-line optimization, however, introduces further sources of approximation, as the gradient w.r.t. the target policy needs to be estimated (off-policy) with samples collected with a behavioral policy. A common choice is to adopt an *importance sampling* (IS) [29, 17] estimator in which each sample is reweighted proportionally to the likelihood of being generated by the target policy. However, directly optimizing this utility function is impractical since it displays a wide variance most of the times [29]. Intuitively, the variance increases proportionally to the distance between the behavioral and the target policy; thus, the estimate is reliable as long as the two policies are close enough. Preventing uncontrolled updates in the space of policy parameters is at the core of the natural gradient approaches [1] applied effectively both on PG methods [18, 33, 63] and on PGPE methods [26]. More recently, this idea has been captured (albeit indirectly) by TRPO, which optimizes via (approximate) natural gradient a surrogate objective function, derived from safe RL [18, 35], subject to a constraint on the Kullback-Leibler divergence between the behavioral and target policy.² Similarly, PPO performs a truncation of the importance weights to discourage the optimization process from going too far. Although TRPO and PPO, together with DDPG, represent the state-of-the-art policy optimization methods in RL for continuous control, they do not explicitly encode in their objective function the uncertainty injected by the importance sampling procedure. A more theoretically grounded analysis has been provided for policy selection [11], model-free [56] and model-based [54] policy evaluation (also accounting for samples collected with multiple behavioral policies), and combined with options [15]. Subsequently, in [55] these methods have been extended for policy improvement, deriving a suitable concentration inequality for the case of truncated importance weights. Unfortunately, these methods are hardly scalable to complex control tasks. A more detailed review of the state-of-the-art policy optimization algorithms is reported in Appendix A.

In this paper, we propose a novel, model-free, actor-only, policy optimization algorithm, named *Policy Optimization via Importance Sampling* (POIS) that mixes online and offline optimization to efficiently exploit the information contained in the collected trajectories. POIS explicitly accounts for the uncertainty introduced by the importance sampling by optimizing a surrogate objective function. The latter captures the trade-off between the estimated performance improvement and the variance injected by the importance sampling. The main contributions of this paper are theoretical, algorithmic and experimental. After revising some notions about importance sampling (Section 3), we propose a concentration inequality, of independent interest, for high-confidence “off-distribution” optimization of objective functions estimated via importance sampling (Section 4). Then we show how this bound can be customized into a surrogate objective function in order to either search in the space of policies (Action-based POIS) or to search in the space of parameters (Parameter-based POIS). The resulting algorithm (in both the action-based and the parameter-based flavor) collects, at each iteration, a set of trajectories. These are used to perform offline optimization of the surrogate objective via gradient ascent (Section 5), after which a new batch of trajectories is collected using the optimized policy. Finally, we provide an experimental evaluation with both linear policies and deep neural policies to illustrate the advantages and limitations of our approach compared to state-of-the-art algorithms (Section 6) on classical control tasks [12, 57]. The proofs for all Theorems and Lemmas can be found in Appendix B. The implementation of POIS can be found at <https://github.com/T3p/pois>.

¹Other solutions to these problems have been proposed in the action-based literature, like the aforementioned DPG algorithm, the gradient baselines [34] and the actor-critic architectures [21].

²Note that this regularization term appears in the performance improvement bound, which contains exact quantities only. Thus, it does not really account for the uncertainty derived from the importance sampling.

2 Preliminaries

A discrete-time Markov Decision Process (MDP) [37] is defined as a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, R, \gamma, D)$ where \mathcal{S} is the state space, \mathcal{A} is the action space, $P(\cdot|s, a)$ is a Markovian transition model that assigns for each state-action pair (s, a) the probability of reaching the next state s' , $\gamma \in [0, 1]$ is the discount factor, $R(s, a) \in [-R_{\max}, R_{\max}]$ assigns the expected reward for performing action a in state s and D is the distribution of the initial state. The behavior of an agent is described by a policy $\pi(\cdot|s)$ that assigns for each state s the probability of performing action a . A trajectory $\tau \in \mathcal{T}$ is a sequence of state-action pairs $\tau = (s_{\tau,0}, a_{\tau,0}, \dots, s_{\tau,H-1}, a_{\tau,H-1}, s_{\tau,H})$, where H is the actual trajectory horizon. The performance of an agent is evaluated in terms of the *expected return*, i.e., the expected discounted sum of the rewards collected along the trajectory: $\mathbb{E}_\tau [R(\tau)]$, where $R(\tau) = \sum_{t=0}^{H-1} \gamma^t R(s_{\tau,t}, a_{\tau,t})$ is the trajectory return.

We focus our attention to the case in which the policy belongs to a parametric policy space $\Pi_\Theta = \{\pi_\theta : \theta \in \Theta \subseteq \mathbb{R}^p\}$. In parameter-based approaches, the agent is equipped with a *hyperpolicy* ν used to sample the policy parameters at the beginning of each episode. The hyperpolicy belongs itself to a parametric hyperpolicy space $\mathcal{N}_\mathcal{P} = \{\nu_\rho : \rho \in \mathcal{P} \subseteq \mathbb{R}^r\}$. The expected return can be expressed, in the parameter-based case, as a double expectation: one over the policy parameter space Θ and one over the trajectory space \mathcal{T} :

$$J_D(\rho) = \int_\Theta \int_{\mathcal{T}} \nu_\rho(\theta) p(\tau|\theta) R(\tau) d\tau d\theta, \quad (1)$$

where $p(\tau|\theta) = D(s_0) \prod_{t=0}^{H-1} \pi_\theta(a_t|s_t) P(s_{t+1}|s_t, a_t)$ is the trajectory density function. The goal of a parameter-based learning agent is to determine the hyperparameters ρ^* so as to maximize $J_D(\rho)$. If ν_ρ is stochastic and differentiable, the hyperparameters can be learned according to the gradient ascent update: $\rho' = \rho + \alpha \nabla_\rho J_D(\rho)$, where $\alpha > 0$ is the step size and $\nabla_\rho J_D(\rho) = \int_\Theta \int_{\mathcal{T}} \nu_\rho(\theta) p(\tau|\theta) \nabla_\rho \log \nu_\rho(\theta) R(\tau) d\tau d\theta$. Since the stochasticity of the hyperpolicy is a sufficient source of exploration, deterministic action policies of the kind $\pi_\theta(a|s) = \delta(a - u_\theta(s))$ are typically considered, where δ is the Dirac delta function and u_θ is a deterministic mapping from \mathcal{S} to \mathcal{A} . In the action-based case, on the contrary, the hyperpolicy ν_ρ is a deterministic distribution $\nu_\rho(\theta) = \delta(\theta - g(\rho))$, where $g(\rho)$ is a deterministic mapping from \mathcal{P} to Θ . For this reason, the dependence on ρ is typically not represented and the expected return expression simplifies into a single expectation over the trajectory space \mathcal{T} :

$$J_D(\theta) = \int_{\mathcal{T}} p(\tau|\theta) R(\tau) d\tau. \quad (2)$$

An action-based learning agent aims to find the policy parameters θ^* that maximize $J_D(\theta)$. In this case, we need to enforce exploration by means of the stochasticity of π_θ . For stochastic and differentiable policies, learning can be performed via gradient ascent: $\theta' = \theta + \alpha \nabla_\theta J_D(\theta)$, where $\nabla_\theta J_D(\theta) = \int_{\mathcal{T}} p(\tau|\theta) \nabla_\theta \log p(\tau|\theta) R(\tau) d\tau$.

3 Evaluation via Importance Sampling

In off-policy evaluation [56, 54], we aim to estimate the performance of a target policy π_T (or hyperpolicy ν_T) given samples collected with a behavioral policy π_B (or hyperpolicy ν_B). More generally, we face the problem of estimating the expected value of a deterministic bounded function f ($\|f\|_\infty < +\infty$) of random variable x taking values in \mathcal{X} under a target distribution P , after having collected samples from a behavioral distribution Q . The *importance sampling* estimator (IS) [7, 29] corrects the distribution with the *importance weights* (or Radon–Nikodym derivative or likelihood ratio) $w_{P/Q}(x) = p(x)/q(x)$:

$$\hat{\mu}_{P/Q} = \frac{1}{N} \sum_{i=1}^N \frac{p(x_i)}{q(x_i)} f(x_i) = \frac{1}{N} \sum_{i=1}^N w_{P/Q}(x_i) f(x_i), \quad (3)$$

where $\mathbf{x} = (x_1, x_2, \dots, x_N)^T$ is sampled from Q and we assume $q(x) > 0$ whenever $f(x)p(x) \neq 0$. This estimator is unbiased ($\mathbb{E}_{\mathbf{x} \sim Q}[\hat{\mu}_{P/Q}] = \mathbb{E}_{x \sim P}[f(x)]$) but it may exhibit an undesirable behavior due to the variability of the importance weights, showing, in some cases, infinite variance. Intuitively, the magnitude of the importance weights provides an indication of how much the probability measures P and Q are dissimilar. This notion can be formalized by the Rényi divergence [40, 59], an information-theoretic dissimilarity index between probability measures.

Rényi divergence Let P and Q be two probability measures on a measurable space $(\mathcal{X}, \mathcal{F})$ such that $P \ll Q$ (P is absolutely continuous w.r.t. Q) and Q is σ -finite. Let P and Q admit p and q as Lebesgue probability density functions (p.d.f.), respectively. The α -Rényi divergence is defined as:

$$D_\alpha(P\|Q) = \frac{1}{\alpha - 1} \log \int_{\mathcal{X}} \left(\frac{dP}{dQ} \right)^\alpha dQ = \frac{1}{\alpha - 1} \log \int_{\mathcal{X}} q(x) \left(\frac{p(x)}{q(x)} \right)^\alpha dx, \quad (4)$$

where dP/dQ is the Radon–Nikodym derivative of P w.r.t. Q and $\alpha \in [0, \infty]$. Some remarkable cases are: $\alpha = 1$ when $D_1(P\|Q) = D_{\text{KL}}(P\|Q)$ and $\alpha = \infty$ yielding $D_\infty(P\|Q) = \log \text{ess sup}_{\mathcal{X}} dP/dQ$. Importing the notation from [8], we indicate the exponentiated α -Rényi divergence as $d_\alpha(P\|Q) = \exp(D_\alpha(P\|Q))$. With little abuse of notation, we will replace $D_\alpha(P\|Q)$ with $D_\alpha(p\|q)$ whenever possible within the context.

The Rényi divergence provides a convenient expression for the moments of the importance weights: $\mathbb{E}_{x \sim Q} [w_{P/Q}(x)^\alpha] = d_\alpha(P\|Q)$. Moreover, $\text{Var}_{x \sim Q} [w_{P/Q}(x)] = d_2(P\|Q) - 1$ and $\text{ess sup}_{x \sim Q} w_{P/Q}(x) = d_\infty(P\|Q)$ [8]. To mitigate the variance problem of the IS estimator, we can resort to the *self-normalized importance sampling* estimator (SN) [7]:

$$\tilde{\mu}_{P/Q} = \frac{\sum_{i=1}^N w_{P/Q}(x_i) f(x_i)}{\sum_{i=1}^N w_{P/Q}(x_i)} = \sum_{i=1}^N \tilde{w}_{P/Q}(x_i) f(x_i), \quad (5)$$

where $\tilde{w}_{P/Q}(x) = w_{P/Q}(x) / \sum_{i=1}^N w_{P/Q}(x_i)$ is the self-normalized importance weight. Differently from $\hat{\mu}_{P/Q}$, $\tilde{\mu}_{P/Q}$ is biased but consistent [29] and it typically displays a more desirable behavior because of its smaller variance.³ Given the realization x_1, x_2, \dots, x_N we can interpret the SN estimator as the expected value of f under an approximation of the distribution P made by N deltas, i.e., $\tilde{p}(x) = \sum_{i=1}^N \tilde{w}_{P/Q}(x_i) \delta(x - x_i)$. The problem of assessing the quality of the SN estimator has been extensively studied by the simulation community, producing several diagnostic indexes to indicate when the weights might display problematic behavior [29]. The *effective sample size* (ESS) was introduced in [22] as the number of samples drawn from P so that the variance of the Monte Carlo estimator $\tilde{\mu}_{P/P}$ is approximately equal to the variance of the SN estimator $\tilde{\mu}_{P/Q}$ computed with N samples. Here we report the original definition and its most common estimate:

$$\text{ESS}(P\|Q) = \frac{N}{\text{Var}_{x \sim Q} [w_{P/Q}(x)] + 1} = \frac{N}{d_2(P\|Q)}, \quad \widehat{\text{ESS}}(P\|Q) = \frac{1}{\sum_{i=1}^N \tilde{w}_{P/Q}(x_i)^2}. \quad (6)$$

The ESS has an interesting interpretation: if $d_2(P\|Q) = 1$, i.e., $P = Q$ almost everywhere, then $\text{ESS} = N$ since we are performing Monte Carlo estimation. Otherwise, the ESS decreases as the dissimilarity between the two distributions increases. In the literature, other ESS-like diagnostics have been proposed that also account for the nature of f [24].

4 Optimization via Importance Sampling

The off-policy optimization problem [55] can be formulated as finding the best target policy π_T (or hyperpolicy ν_T), i.e., the one maximizing the expected return, having access to a set of samples collected with a behavioral policy π_B (or hyperpolicy ν_B). In a more abstract sense, we aim to determine the target distribution P that maximizes $\mathbb{E}_{x \sim P}[f(x)]$ having samples collected from the fixed behavioral distribution Q . In this section, we analyze the problem of defining a proper objective function for this purpose. Directly optimizing the estimator $\hat{\mu}_{P/Q}$ or $\tilde{\mu}_{P/Q}$ is, in most of the cases, unsuccessful. With enough freedom in choosing P , the optimal solution would assign as much probability mass as possible to the maximum value among $f(x_i)$. Clearly, in this scenario, the estimator is unreliable and displays a large variance. For this reason, we adopt a risk-averse approach and we decide to optimize a statistical *lower bound* of the expected value $\mathbb{E}_{x \sim P}[f(x)]$ that holds with high confidence. We start by analyzing the behavior of the IS estimator and we provide the following result that bounds the variance of $\hat{\mu}_{P/Q}$ in terms of the Rényi divergence.

Lemma 4.1. *Let P and Q be two probability measures on the measurable space $(\mathcal{X}, \mathcal{F})$ such that $P \ll Q$. Let $\mathbf{x} = (x_1, x_2, \dots, x_N)^T$ i.i.d. random variables sampled from Q and $f : \mathcal{X} \rightarrow \mathbb{R}$ be a*

³Note that $|\tilde{\mu}_{P/Q}| \leq \|f\|_\infty$. Therefore, its variance is always finite.

bounded function ($\|f\|_\infty < +\infty$). Then, for any $N > 0$, the variance of the IS estimator $\hat{\mu}_{P/Q}$ can be upper bounded as:

$$\text{Var}_{\mathbf{x} \sim Q} [\hat{\mu}_{P/Q}] \leq \frac{1}{N} \|f\|_\infty^2 d_2(P\|Q). \quad (7)$$

When $P = Q$ almost everywhere, we get $\text{Var}_{\mathbf{x} \sim Q} [\hat{\mu}_{Q/Q}] \leq \frac{1}{N} \|f\|_\infty^2$, a well-known bound on the variance of a Monte Carlo estimator. Recalling the definition of ESS (6) we can rewrite the previous bound as: $\text{Var}_{\mathbf{x} \sim Q} [\hat{\mu}_{P/Q}] \leq \frac{\|f\|_\infty^2}{\text{ESS}(P\|Q)}$, i.e., the variance scales with ESS instead of N . While $\hat{\mu}_{P/Q}$ can have unbounded variance even if f is bounded, the SN estimator $\tilde{\mu}_{P/Q}$ is always bounded by $\|f\|_\infty$ and therefore it always has a finite variance. Since the normalization term makes all the samples $\tilde{w}_{P/Q}(x_i)f(x_i)$ interdependent, an exact analysis of its bias and variance is more challenging. Several works adopted approximate methods to provide an expression for the variance [17]. We propose an analysis of bias and variance of the SN estimator in Appendix D.

4.1 Concentration Inequality

Finding a suitable concentration inequality for off-policy learning was studied in [56] for offline policy evaluation and subsequently in [55] for optimization. On one hand, fully empirical concentration inequalities, like Student-T, besides the asymptotic approximation, are not suitable in this case since the empirical variance needs to be estimated with importance sampling as well injecting further uncertainty [29]. On the other hand, several distribution-free inequalities like Hoeffding require knowing the maximum of the estimator, which might not exist ($d_\infty(P\|Q) = \infty$) for the IS estimator. Constraining $d_\infty(P\|Q)$ to be finite often introduces unacceptable limitations. For instance, in the case of univariate Gaussian distributions, it prevents a step that selects a target variance larger than the behavioral one from being performed (see Appendix C).⁴ Even Bernstein inequalities [5], are hardly applicable since, for instance, in the case of univariate Gaussian distributions, the importance weights display a fat tail behavior (see Appendix C). We believe that a reasonable trade-off is to require the variance of the importance weights to be finite, that is equivalent to require $d_2(P\|Q) < \infty$, i.e., $\sigma_P < 2\sigma_Q$ for univariate Gaussians. For this reason, we resort to Chebyshev-like inequalities and we propose the following concentration bound derived from Cantelli's inequality and customized for the IS estimator.

Theorem 4.1. *Let P and Q be two probability measures on the measurable space $(\mathcal{X}, \mathcal{F})$ such that $P \ll Q$ and $d_2(P\|Q) < +\infty$. Let x_1, x_2, \dots, x_N be i.i.d. random variables sampled from Q , and $f : \mathcal{X} \rightarrow \mathbb{R}$ be a bounded function ($\|f\|_\infty < +\infty$). Then, for any $0 < \delta \leq 1$ and $N > 0$ with probability at least $1 - \delta$ it holds that:*

$$\mathbb{E}_{\mathbf{x} \sim P} [f(x)] \geq \frac{1}{N} \sum_{i=1}^N w_{P/Q}(x_i) f(x_i) - \|f\|_\infty \sqrt{\frac{(1-\delta)d_2(P\|Q)}{\delta N}}. \quad (8)$$

The bound highlights the interesting trade-off between the estimated performance and the uncertainty introduced by changing the distribution. The latter enters in the bound as the 2-Rényi divergence between the target distribution P and the behavioral distribution Q . Intuitively, we should trust the estimator $\hat{\mu}_{P/Q}$ as long as P is not too far from Q . For the SN estimator, accounting for the bias, we are able to obtain a bound (reported in Appendix D), with a similar dependence on P as in Theorem 4.1, albeit with different constants. Renaming all constants involved in the bound of Theorem 4.1 as $\lambda = \|f\|_\infty \sqrt{(1-\delta)/\delta}$, we get a surrogate objective function. The optimization can be carried out in different ways. The following section shows why using the natural gradient could be a successful choice in case P and Q can be expressed as parametric differentiable distributions.

4.2 Importance Sampling and Natural Gradient

We can look at a parametric distribution P_ω , having p_ω as a density function, as a point on a probability manifold with coordinates $\omega \in \Omega$. If p_ω is differentiable, the Fisher Information Matrix (FIM) [39, 2] is defined as: $\mathcal{F}(\omega) = \int_{\mathcal{X}} p_\omega(x) \nabla_\omega \log p_\omega(x) \nabla_\omega \log p_\omega(x)^T dx$. This matrix is, up to a scale, an

⁴Although the variance tends to be reduced in the learning process, there might be cases in which it needs to be increased (e.g., suppose we start with a behavioral policy with small variance, it might be beneficial increasing the variance to enforce exploration).

Algorithm 1 Action-based POIS

Initialize θ_0^0 arbitrarily
for $j = 0, 1, 2, \dots$, until convergence **do**
 Collect N trajectories with $\pi_{\theta_0^j}$
 for $k = 0, 1, 2, \dots$, until convergence **do**
 Compute $\mathcal{G}(\theta_k^j)$, $\nabla_{\theta_k^j} \mathcal{L}(\theta_k^j / \theta_0^j)$ and α_k
 $\theta_{k+1}^j = \theta_k^j + \alpha_k \mathcal{G}(\theta_k^j)^{-1} \nabla_{\theta_k^j} \mathcal{L}(\theta_k^j / \theta_0^j)$
 end for
 $\theta_0^{j+1} = \theta_k^j$
end for

Algorithm 2 Parameter-based POIS

Initialize ρ_0^0 arbitrarily
for $j = 0, 1, 2, \dots$, until convergence **do**
 Sample N policy parameters θ_i^j from $\nu_{\rho_0^j}$
 Collect a trajectory with each $\pi_{\theta_i^j}$
 for $k = 0, 1, 2, \dots$, until convergence **do**
 Compute $\mathcal{G}(\rho_k^j)$, $\nabla_{\rho_k^j} \mathcal{L}(\rho_k^j / \rho_0^j)$ and α_k
 $\rho_{k+1}^j = \rho_k^j + \alpha_k \mathcal{G}(\rho_k^j)^{-1} \nabla_{\rho_k^j} \mathcal{L}(\rho_k^j / \rho_0^j)$
 end for
 $\rho_0^{j+1} = \rho_k^j$
end for

invariant metric [1] on parameter space Ω , i.e., $\kappa(\omega' - \omega)^T \mathcal{F}(\omega)(\omega' - \omega)$ is independent on the specific parameterization and provides a second order approximation of the distance between p_ω and $p_{\omega'}$ on the probability manifold up to a scale factor $\kappa \in \mathbb{R}$. Given a loss function $\mathcal{L}(\omega)$, we define the natural gradient [1, 19] as $\tilde{\nabla}_\omega \mathcal{L}(\omega) = \mathcal{F}^{-1}(\omega) \nabla_\omega \mathcal{L}(\omega)$, which represents the steepest ascent direction in the probability manifold. Thanks to the invariance property, there is a tight connection between the geometry induced by the Rényi divergence and the Fisher information metric [3].

Theorem 4.2. *Let p_ω be a p.d.f. differentiable w.r.t. $\omega \in \Omega$. Then, it holds that, for the Rényi divergence: $D_\alpha(p_{\omega'} \| p_\omega) = \frac{\alpha}{2} (\omega' - \omega)^T \mathcal{F}(\omega) (\omega' - \omega) + o(\|\omega' - \omega\|_2^2)$, and for the exponentiated Rényi divergence: $d_\alpha(p_{\omega'} \| p_\omega) = 1 + \frac{\alpha}{2} (\omega' - \omega)^T \mathcal{F}(\omega) (\omega' - \omega) + o(\|\omega' - \omega\|_2^2)$.*

This result provides an approximate expression for the variance of the importance weights, as $\text{Var}_{x \sim p_\omega} [w_{\omega' / \omega}(x)] = d_2(p_{\omega'} \| p_\omega) - 1 \simeq \frac{\alpha}{2} (\omega' - \omega)^T \mathcal{F}(\omega) (\omega' - \omega)$. It also justifies the use of natural gradients in off-distribution optimization, since a step in natural gradient direction has a controllable effect on the variance of the importance weights.

5 Policy Optimization via Importance Sampling

In this section, we discuss how to customize the bound provided in Theorem 4.1 for policy optimization, developing a novel model-free actor-only policy search algorithm, named *Policy Optimization via Importance Sampling* (POIS). We propose two versions of POIS: *Action-based POIS* (A-POIS), which is based on a policy gradient approach, and *Parameter-based POIS* (P-POIS), which adopts the PGPE framework. A more detailed description of the implementation aspects is reported in Appendix E.

5.1 Action-based POIS

In Action-based POIS (A-POIS) we search for a policy that maximizes the performance index $J_D(\theta)$ within a parametric space $\Pi_\Theta = \{\pi_\theta : \theta \in \Theta \subseteq \mathbb{R}^p\}$ of stochastic differentiable policies. In this context, the behavioral (resp. target) distribution Q (resp. P) becomes the distribution over trajectories $p(\cdot | \theta)$ (resp. $p(\cdot | \theta')$) induced by the behavioral policy π_θ (resp. target policy $\pi_{\theta'}$) and f is the trajectory return $R(\tau)$ which is uniformly bounded as $|R(\tau)| \leq R_{\max} \frac{1-\gamma^H}{1-\gamma}$.⁵ The surrogate loss function cannot be directly optimized via gradient ascent since computing $d_\alpha(p(\cdot | \theta') \| p(\cdot | \theta))$ requires the approximation of an integral over the trajectory space and, for stochastic environments, to know the transition model P , which is unknown in a model-free setting. Simple bounds to this quantity, like $d_\alpha(p(\cdot | \theta') \| p(\cdot | \theta)) \leq \sup_{s \in \mathcal{S}} d_\alpha(\pi_{\theta'}(\cdot | s) \| \pi_\theta(\cdot | s))^H$, besides being hard to compute due to the presence of the supremum, are extremely conservative since the Rényi divergence is raised to the horizon H . We suggest the replacement of the Rényi divergence with an estimate $\hat{d}_2(p(\cdot | \theta') \| p(\cdot | \theta)) = \frac{1}{N} \sum_{i=1}^N \prod_{t=0}^{H-1} d_2(\pi_{\theta'}(\cdot | s_{\tau_i, t}) \| \pi_\theta(\cdot | s_{\tau_i, t}))$ defined only in terms of the policy Rényi divergence (see Appendix E.2 for details). Thus, we obtain the following surrogate

⁵When $\gamma \rightarrow 1$ the bound becomes HR_{\max} .

objective:

$$\mathcal{L}_\lambda^{\text{A-POIS}}(\theta'/\theta) = \frac{1}{N} \sum_{i=1}^N w_{\theta'/\theta}(\tau_i) R(\tau_i) - \lambda \sqrt{\frac{\widehat{d}_2(p(\cdot|\theta')\|p(\cdot|\theta))}{N}}, \quad (9)$$

where $w_{\theta'/\theta}(\tau_i) = \frac{p(\tau_i|\theta')}{p(\tau_i|\theta)} = \prod_{t=0}^{H-1} \frac{\pi_{\theta'}(a_{\tau_{i,t}}|s_{\tau_{i,t}})}{\pi_{\theta}(a_{\tau_{i,t}}|s_{\tau_{i,t}})}$. We consider the case in which $\pi_{\theta}(\cdot|s)$ is a Gaussian distribution over actions whose mean depends on the state and whose covariance is state-independent and diagonal: $\mathcal{N}(u_{\mu}(s), \text{diag}(\sigma^2))$, where $\theta = (\mu, \sigma)$. The learning process mixes online and offline optimization. At each online iteration j , a dataset of N trajectories is collected by executing in the environment the current policy $\pi_{\theta_0^j}$. These trajectories are used to optimize the surrogate loss function $\mathcal{L}_\lambda^{\text{A-POIS}}$. At each offline iteration k , the parameters are updated via gradient ascent: $\theta_{k+1}^j = \theta_k^j + \alpha_k \mathcal{G}(\theta_k^j)^{-1} \nabla_{\theta_k^j} \mathcal{L}(\theta_k^j/\theta_0^j)$, where $\alpha_k > 0$ is the step size which is chosen via line search (see Appendix E.1) and $\mathcal{G}(\theta_k^j)$ is a positive semi-definite matrix (e.g., $\mathcal{F}(\theta_k^j)$, the FIM, for natural gradient)⁶. The pseudo-code of POIS is reported in Algorithm 1.

5.2 Parameter-based POIS

In the Parameter-based POIS (P-POIS) we again consider a parametrized policy space $\Pi_{\Theta} = \{\pi_{\theta} : \theta \in \Theta \subseteq \mathbb{R}^p\}$, but π_{θ} needs not be differentiable. The policy parameters θ are sampled at the beginning of each episode from a parametric hyperpolicy ν_{ρ} selected in a parametric space $\mathcal{N}_{\mathcal{P}} = \{\nu_{\rho} : \rho \in \mathcal{P} \subseteq \mathbb{R}^r\}$. The goal is to learn the *hyperparameters* ρ so as to maximize $J_D(\rho)$. In this setting, the distributions Q and P of Section 4 correspond to the behavioral ν_{ρ} and target $\nu_{\rho'}$ hyperpolicies, while f remains the trajectory return $R(\tau)$. The importance weights [67] must take into account all sources of randomness, derived from sampling a policy parameter θ and a trajectory τ : $w_{\rho'/\rho}(\theta) = \frac{\nu_{\rho'}(\theta)p(\tau|\theta)}{\nu_{\rho}(\theta)p(\tau|\theta)} = \frac{\nu_{\rho'}(\theta)}{\nu_{\rho}(\theta)}$. In practice, a Gaussian hyperpolicy ν_{ρ} with diagonal covariance matrix is often used, i.e., $\mathcal{N}(\mu, \text{diag}(\sigma^2))$ with $\rho = (\mu, \sigma)$. The policy is assumed to be deterministic: $\pi_{\theta}(a|s) = \delta(a - u_{\theta}(s))$, where u_{θ} is a deterministic function of the state s [e.g., 46, 14]. A first advantage over the action-based setting is that the distribution of the importance weights is entirely known, as it is the ratio of two Gaussians and the Rényi divergence $d_2(\nu_{\rho'}\|\nu_{\rho})$ can be computed exactly [6] (see Appendix C). This leads to the following surrogate objective:

$$\mathcal{L}_\lambda^{\text{P-POIS}}(\rho'/\rho) = \frac{1}{N} \sum_{i=1}^N w_{\rho'/\rho}(\theta_i) R(\tau_i) - \lambda \sqrt{\frac{d_2(\nu_{\rho'}\|\nu_{\rho})}{N}}, \quad (10)$$

where each trajectory τ_i is obtained by running an episode with action policy π_{θ_i} , and the corresponding policy parameters θ_i are sampled independently from hyperpolicy ν_{ρ} , at the beginning of each episode. The hyperpolicy parameters are then updated offline as $\rho_{k+1}^j = \rho_k^j + \alpha_k \mathcal{G}(\rho_k^j)^{-1} \nabla_{\rho_k^j} \mathcal{L}(\rho_k^j/\rho_0^j)$ (see Algorithm 2 for the complete pseudo-code). A further advantage w.r.t. the action-based case is that the FIM $\mathcal{F}(\rho)$ can be computed exactly, and it is diagonal in the case of a Gaussian hyperpolicy with diagonal covariance matrix, turning a problematic inversion into a trivial division (the FIM is block-diagonal in the more general case of a Gaussian hyperpolicy, as observed in [26]). This makes natural gradient much more enticing for P-POIS.

6 Experimental Evaluation

In this section, we present the experimental evaluation of POIS in its two flavors (action-based and parameter-based). We first provide a set of empirical comparisons on classical continuous control tasks with linearly parametrized policies; we then show how POIS can be also adopted for learning deep neural policies. In all experiments, for the A-POIS we used the IS estimator, while for P-POIS we employed the SN estimator. All experimental details are provided in Appendix F.

6.1 Linear Policies

Linear parametrized Gaussian policies proved their ability to scale on complex control tasks [38]. In this section, we compare the learning performance of A-POIS and P-POIS against TRPO [42] and

⁶The FIM needs to be estimated via importance sampling as well, as shown in Appendix E.3.

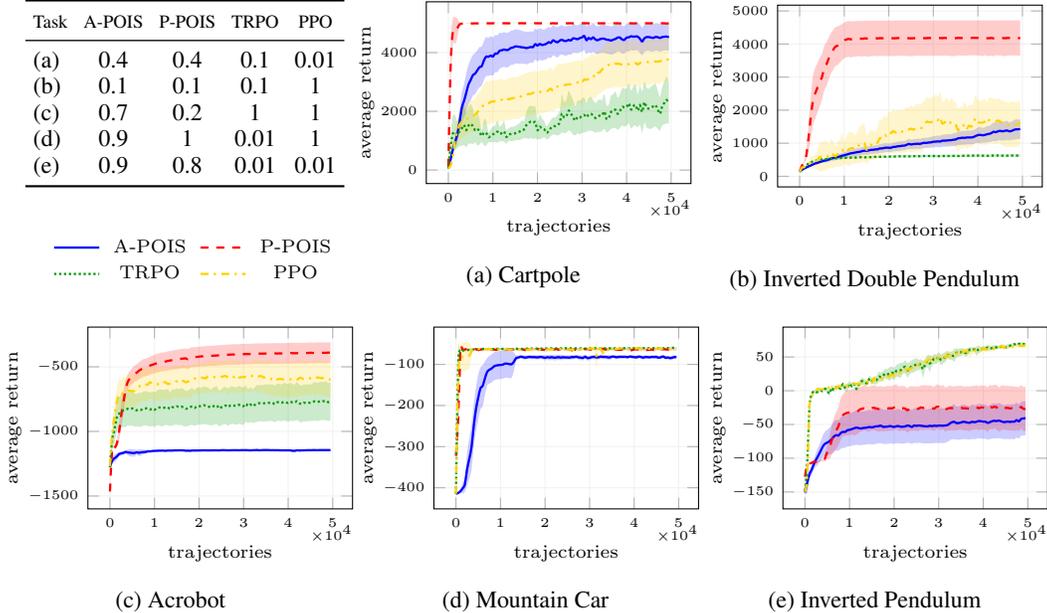


Figure 1: Average return as a function of the number of trajectories for A-POIS, P-POIS and TRPO with *linear policy* (20 runs, 95% c.i.). The table reports the best hyperparameters found (δ for POIS and the step size for TRPO and PPO).

PPO [44] on classical continuous control benchmarks [12]. In Figure 1, we can see that both versions of POIS are able to significantly outperform both TRPO and PPO in the Cartpole environments, especially the P-POIS. In the Inverted Double Pendulum environment the learning curve of P-POIS is remarkable while A-POIS displays a behavior comparable to PPO. In the Acrobot task, P-POIS displays a better performance w.r.t. TRPO and PPO, but A-POIS does not keep up. In Mountain Car, we see yet another behavior: the learning curves of TRPO, PPO and P-POIS are almost one-shot (even if PPO shows a small instability), while A-POIS fails to display such a fast convergence. Finally, in the Inverted Pendulum environment, TRPO and PPO outperform both versions of POIS. This example highlights a limitation of our approach. Since POIS performs an importance sampling procedure at trajectory level, it cannot assign credit to good actions in bad trajectories. On the contrary, weighting each sample, TRPO and PPO are able also to exploit good trajectory segments. In principle, this problem can be mitigated in POIS by resorting to *per-decision importance sampling* [36], in which the weight is assigned to individual rewards instead of trajectory returns. Overall, POIS displays a performance comparable with TRPO and PPO across the tasks. In particular, P-POIS displays a better performance w.r.t. A-POIS. However, this ordering is not maintained when moving to more complex policy architectures, as shown in the next section.

In Figure 2 we show, for several metrics, the behavior of A-POIS when changing the δ parameter in the Cartpole environment. We can see that when δ is small (e.g., 0.2), the Effective Sample Size (ESS) remains large and, consequently, the variance of the importance weights ($\text{Var}[w]$) is small. This means that the penalization term in the objective function discourages the optimization process from selecting policies which are far from the behavioral policy. As a consequence, the displayed behavior is very conservative, preventing the policy from reaching the optimum. On the contrary, when δ approaches 1, the ESS is smaller and the variance of the weights tends to increase significantly. Again, the performance remains suboptimal as the penalization term in the objective function is too light. The best behavior is obtained with an intermediate value of δ , specifically 0.4.

6.2 Deep Neural Policies

In this section, we adopt a deep neural network (3 layers: 100, 50, 25 neurons each) to represent the policy. The experiment setup is fully compatible with the classical benchmark [12]. While A-POIS can be directly applied to deep neural networks, P-POIS exhibits some critical issues. A highly dimensional hyperpolicy (like a Gaussian from which the weights of an MLP policy are

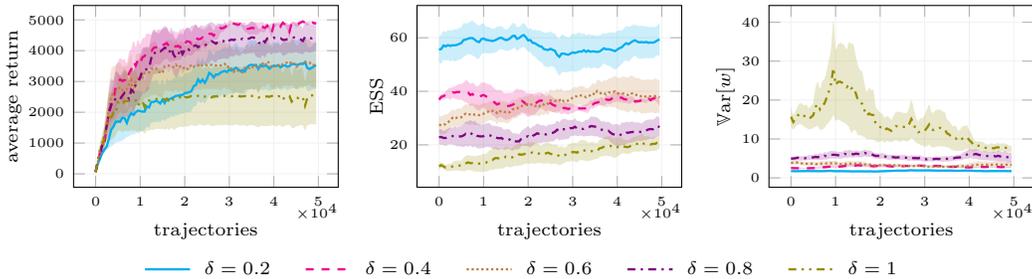


Figure 2: Average return, Effective Sample Size (ESS) and variance of the importance weights ($\text{Var}[w]$) as a function of the number of trajectories for A-POIS for different values of the parameter δ in the Cartpole environment (20 runs, 95% c.i.).

Table 1: Performance of POIS compared with [12] on *deep neural policies* (5 runs, 95% c.i.). In **bold**, the performances that are not statistically significantly different from the best algorithm in each task.

Algorithm	Cart-Pole Balancing	Mountain Car	Double Inverted Pendulum	Swimmer
REINFORCE	4693.7 \pm 14.0	-67.1 \pm 1.0	4116.5 \pm 65.2	92.3 \pm 0.1
TRPO	4869.8 \pm 37.6	-61.7 \pm 0.9	4412.4 \pm 50.4	96.0 \pm 0.2
DDPG	4634.4 \pm 87.6	-288.4 \pm 170.3	2863.4 \pm 154.0	85.8 \pm 1.8
A-POIS	4842.8 \pm 13.0	-63.7 \pm 0.5	4232.1 \pm 189.5	88.7 \pm 0.55
CEM	4815.4 \pm 4.8	-66.0 \pm 2.4	2566.2 \pm 178.9	68.8 \pm 2.4
P-POIS	4428.1 \pm 138.6	-78.9 \pm 2.5	3161.4 \pm 959.2	76.8 \pm 1.6

sampled) can make $d_2(\nu_{\rho'} \parallel \nu_{\rho})$ extremely sensitive to small parameter changes, leading to over-conservative updates.⁷ A first practical variant comes from the insight that $d_2(\nu_{\rho'} \parallel \nu_{\rho})/N$ is the inverse of the effective sample size, as reported in Equation 6. We can obtain a less conservative (although approximate) surrogate function by replacing it with $1/\widehat{\text{ESS}}(\nu_{\rho'} \parallel \nu_{\rho})$. Another trick is to model the hyperpolicy as a set of independent Gaussians, each defined over a disjoint subspace of Θ (implementation details are provided in Appendix E.5). In Table 1, we augmented the results provided in [12] with the performance of POIS for the considered tasks. We can see that A-POIS is able to reach an overall behavior comparable with the best of the action-based algorithms, approaching TRPO and beating DDPG. Similarly, P-POIS exhibits a performance similar to CEM [52], the best performing among the parameter-based methods. The complete results are reported in Appendix F.

7 Discussion and Conclusions

In this paper, we presented a new actor-only policy optimization algorithm, POIS, which alternates online and offline optimization in order to efficiently exploit the collected trajectories, and can be used in combination with action-based and parameter-based exploration. In contrast to the state-of-the-art algorithms, POIS has a strong theoretical grounding, since its surrogate objective function derives from a statistical bound on the estimated performance, that is able to capture the uncertainty induced by importance sampling. The experimental evaluation showed that POIS, in both its versions (action-based and parameter-based), is able to achieve a performance comparable with TRPO, PPO and other classical algorithms on continuous control tasks. Natural extensions of POIS could focus on employing per-decision importance sampling, adaptive batch size, and trajectory reuse. Future work also includes scaling POIS to high-dimensional tasks and highly-stochastic environments. We believe that this work represents a valuable starting point for a deeper understanding of modern policy optimization and for the development of effective and scalable policy search methods.

⁷This curse of dimensionality, related to $\dim(\theta)$, has some similarities with the dependence of the Rényi divergence on the actual horizon H in the action-based case.

Acknowledgments

The study was partially funded by Lombardy Region (Announcement PORFESR 2014-2020). F. F. was partially funded through ERC Advanced Grant (no: 742870). We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Tesla K40cm, Titan XP and Tesla V100 used for this research.

References

- [1] Shun-Ichi Amari. Natural gradient works efficiently in learning. *Neural computation*, 10(2):251–276, 1998.
- [2] Shun-ichi Amari. *Differential-geometrical methods in statistics*, volume 28. Springer Science & Business Media, 2012.
- [3] Shun-ichi Amari and Andrzej Cichocki. Information geometry of divergence functions. *Bulletin of the Polish Academy of Sciences: Technical Sciences*, 58(1):183–195, 2010.
- [4] Jonathan Baxter and Peter L Bartlett. Infinite-horizon policy-gradient estimation. *Journal of Artificial Intelligence Research*, 15:319–350, 2001.
- [5] Bernard Bercu, Bernard Delyon, and Emmanuel Rio. Concentration inequalities for sums. In *Concentration Inequalities for Sums and Martingales*, pages 11–60. Springer, 2015.
- [6] Jacob Burbea. The convexity with respect to gaussian distributions of divergences of order α . *Utilitas Mathematica*, 26:171–192, 1984.
- [7] William G Cochran. *Sampling techniques*. John Wiley & Sons, 2007.
- [8] Corinna Cortes, Yishay Mansour, and Mehryar Mohri. Learning bounds for importance weighting. In *Advances in neural information processing systems*, pages 442–450, 2010.
- [9] Thomas Degris, Martha White, and Richard S Sutton. Off-policy actor-critic. *arXiv preprint arXiv:1205.4839*, 2012.
- [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] Shayan Doroudi, Philip S Thomas, and Emma Brunskill. Importance sampling for fair policy selection. *Uncertainty in Artificial Intelligence*, 2017.
- [12] Yan Duan, Xi Chen, Rein Houthoofd, John Schulman, and Pieter Abbeel. Benchmarking deep reinforcement learning for continuous control. In *International Conference on Machine Learning*, pages 1329–1338, 2016.
- [13] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256, 2010.
- [14] Mandy Grüttner, Frank Sehnke, Tom Schaul, and Jürgen Schmidhuber. Multi-dimensional deep memory go-player for parameter exploring policy gradients. 2010.
- [15] Zhaohan Guo, Philip S Thomas, and Emma Brunskill. Using options and covariance testing for long horizon off-policy policy evaluation. In *Advances in Neural Information Processing Systems*, pages 2489–2498, 2017.
- [16] Nikolaus Hansen and Andreas Ostermeier. Completely derandomized self-adaptation in evolution strategies. *Evolutionary computation*, 9(2):159–195, 2001.
- [17] Timothy Classen Hesterberg. *Advances in importance sampling*. PhD thesis, Stanford University, 1988.
- [18] Sham Kakade and John Langford. Approximately optimal approximate reinforcement learning. In *International Conference on Machine Learning*, volume 2, pages 267–274, 2002.
- [19] Sham M Kakade. A natural policy gradient. In *Advances in neural information processing systems*, pages 1531–1538, 2002.
- [20] Jens Kober, J Andrew Bagnell, and Jan Peters. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11):1238–1274, 2013.

- [21] Vijay R Konda and John N Tsitsiklis. Actor-critic algorithms. In *Advances in neural information processing systems*, pages 1008–1014, 2000.
- [22] Augustine Kong. A note on importance sampling using standardized weights. *University of Chicago, Dept. of Statistics, Tech. Rep.*, 348, 1992.
- [23] 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.
- [24] Luca Martino, Víctor Elvira, and Francisco Louzada. Effective sample size for importance sampling based on discrepancy measures. *Signal Processing*, 131:386–401, 2017.
- [25] Takamitsu Matsuura, Tetsuro Morimura, and Jun Morimoto. Adaptive step-size policy gradients with average reward metric. In *Proceedings of 2nd Asian Conference on Machine Learning*, pages 285–298, 2010.
- [26] Atsushi Miyamae, Yuichi Nagata, Isao Ono, and Shigenobu Kobayashi. Natural policy gradient methods with parameter-based exploration for control tasks. In *Advances in neural information processing systems*, pages 1660–1668, 2010.
- [27] Rémi Munos, Tom Stepleton, Anna Harutyunyan, and Marc Bellemare. Safe and efficient off-policy reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 1054–1062, 2016.
- [28] Andrew Y Ng and Michael Jordan. Pegasus: A policy search method for large mdps and pomdps. In *Proceedings of the Sixteenth conference on Uncertainty in artificial intelligence*, pages 406–415. Morgan Kaufmann Publishers Inc., 2000.
- [29] Art B. Owen. *Monte Carlo theory, methods and examples*. 2013.
- [30] Jing Peng and Ronald J Williams. Incremental multi-step q-learning. In *Machine Learning Proceedings 1994*, pages 226–232. Elsevier, 1994.
- [31] Jan Peters, Katharina Mülling, and Yasemin Altun. Relative entropy policy search. In *AAAI*, pages 1607–1612. Atlanta, 2010.
- [32] Jan Peters and Stefan Schaal. Reinforcement learning by reward-weighted regression for operational space control. In *Proceedings of the 24th international conference on Machine learning*, pages 745–750. ACM, 2007.
- [33] Jan Peters and Stefan Schaal. Natural actor-critic. *Neurocomputing*, 71(7-9):1180–1190, 2008.
- [34] Jan Peters and Stefan Schaal. Reinforcement learning of motor skills with policy gradients. *Neural networks*, 21(4):682–697, 2008.
- [35] Matteo Pirota, Marcello Restelli, Alessio Pecorino, and Daniele Calandriello. Safe policy iteration. In *International Conference on Machine Learning*, pages 307–315, 2013.
- [36] Doina Precup, Richard S Sutton, and Satinder P Singh. Eligibility traces for off-policy policy evaluation. In *International Conference on Machine Learning*, pages 759–766. Citeseer, 2000.
- [37] Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- [38] Aravind Rajeswaran, Kendall Lowrey, Emanuel V Todorov, and Sham M Kakade. Towards generalization and simplicity in continuous control. In *Advances in Neural Information Processing Systems*, pages 6553–6564, 2017.
- [39] C Radhakrishna Rao. Information and the accuracy attainable in the estimation of statistical parameters. In *Breakthroughs in statistics*, pages 235–247. Springer, 1992.
- [40] Alfréd Rényi. On measures of entropy and information. Technical report, Hungarian Academy of Sciences Budapest Hungary, 1961.
- [41] Reuven Rubinfeld. The cross-entropy method for combinatorial and continuous optimization. *Methodology and computing in applied probability*, 1(2):127–190, 1999.
- [42] 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.

- [43] John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*, 2015.
- [44] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [45] Frank Sehnke, Christian Osendorfer, Thomas Rückstieß, Alex Graves, Jan Peters, and Jürgen Schmidhuber. Policy gradients with parameter-based exploration for control. In *International Conference on Artificial Neural Networks*, pages 387–396. Springer, 2008.
- [46] Frank Sehnke, Christian Osendorfer, Thomas Rückstieß, Alex Graves, Jan Peters, and Jürgen Schmidhuber. Parameter-exploring policy gradients. *Neural Networks*, 23(4):551–559, 2010.
- [47] David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *International Conference on Machine Learning*, 2014.
- [48] Kenneth O Stanley and Risto Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2):99–127, 2002.
- [49] Yi Sun, Daan Wierstra, Tom Schaul, and Juergen Schmidhuber. Efficient natural evolution strategies. In *Proceedings of the 11th Annual conference on Genetic and evolutionary computation*, pages 539–546. ACM, 2009.
- [50] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.
- [51] 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.
- [52] István Szita and András Lörincz. Learning tetris using the noisy cross-entropy method. *Neural computation*, 18(12):2936–2941, 2006.
- [53] Russ Tedrake, Teresa Weirui Zhang, and H Sebastian Seung. Stochastic policy gradient reinforcement learning on a simple 3d biped. In *Intelligent Robots and Systems, 2004.(IROS 2004). Proceedings. 2004 IEEE/RSJ International Conference on*, volume 3, pages 2849–2854. IEEE, 2004.
- [54] Philip Thomas and Emma Brunskill. Data-efficient off-policy policy evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 2139–2148, 2016.
- [55] Philip Thomas, Georgios Theodorou, and Mohammad Ghavamzadeh. High confidence policy improvement. In *International Conference on Machine Learning*, pages 2380–2388, 2015.
- [56] Philip S Thomas, Georgios Theodorou, and Mohammad Ghavamzadeh. High-confidence off-policy evaluation. In *AAAI*, pages 3000–3006, 2015.
- [57] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pages 5026–5033. IEEE, 2012.
- [58] George Tucker, Surya Bhupatiraju, Shixiang Gu, Richard E Turner, Zoubin Ghahramani, and Sergey Levine. The mirage of action-dependent baselines in reinforcement learning. *arXiv preprint arXiv:1802.10031*, 2018.
- [59] Tim Van Erven and Peter Harremoës. Rényi divergence and kullback-leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014.
- [60] Jay M Ver Hoef. Who invented the delta method? *The American Statistician*, 66(2):124–127, 2012.
- [61] Ziyu Wang, Victor Bapst, Nicolas Heess, Volodymyr Mnih, Remi Munos, Koray Kavukcuoglu, and Nando de Freitas. Sample efficient actor-critic with experience replay. *arXiv preprint arXiv:1611.01224*, 2016.
- [62] Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.
- [63] Daan Wierstra, Tom Schaul, Jan Peters, and Juergen Schmidhuber. Natural evolution strategies. In *Evolutionary Computation, 2008. CEC 2008.(IEEE World Congress on Computational Intelligence). IEEE Congress on*, pages 3381–3387. IEEE, 2008.
- [64] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. In *Reinforcement Learning*, pages 5–32. Springer, 1992.

- [65] Cathy Wu, Aravind Rajeswaran, Yan Duan, Vikash Kumar, Alexandre M Bayen, Sham Kakade, Igor Mordatch, and Pieter Abbeel. Variance reduction for policy gradient with action-dependent factorized baselines. *arXiv preprint arXiv:1803.07246*, 2018.
- [66] Tingting Zhao, Hirotaka Hachiya, Gang Niu, and Masashi Sugiyama. Analysis and improvement of policy gradient estimation. In *Advances in Neural Information Processing Systems*, pages 262–270, 2011.
- [67] Tingting Zhao, Hirotaka Hachiya, Voot Tangkaratt, Jun Morimoto, and Masashi Sugiyama. Efficient sample reuse in policy gradients with parameter-based exploration. *Neural computation*, 25(6):1512–1547, 2013.

Index of the Appendix

In the following, we briefly recap the contents of the Appendix.

- Appendix A shows a more detailed comparison of POIS with the policy-search algorithms. Table 2 summarizes some features of the considered methods.
- Appendix B reports all proofs and derivations.
- Appendix C provides an analysis of the distribution of the importance weights in the case of univariate Gaussian behavioral and target distributions.
- Appendix D shows some bounds on bias and variance for the self-normalized importance sampling estimator and provides a high confidence bound.
- Appendix E illustrates some implementation details of POIS, in particular line search algorithms, estimation of the Rényi divergence, computation of the FIM and practical versions of P-POIS.
- Appendix F provides the hyperparameters used in the experiments and further results.

A Related Works

Policy optimization algorithms can be classified according to different dimensions (Table 2). It is by now established, in the policy-based RL community, that effective algorithms, either on-policy or off-policy, should account for the variance of the gradient estimate. Early attempts, in the class of action-based algorithms, are the usage of a baseline to reduce the estimated gradient variance without introducing bias [4, 34]. A similar rationale is at the basis of actor-critic architectures [21, 51, 33], in which an estimate of the value function is used to reduce uncertainty. Baselines are typically constant (REINFORCE), time-dependent (G(PO)MDP) or state-dependent (actor-critic), but these approaches have been recently extended to account for action-dependent baselines [58, 65]. Even though parameter-based algorithms are, by nature, affected by smaller variance w.r.t. action-based ones, similar baselines can be derived [66]. A first dichotomy in the class of policy-based algorithms comes when considering the minimal unit used to compute the gradient. *Episode-based* (or episodic) approaches [e.g., 64, 4] perform the gradient estimation by averaging the gradients of each episode which need to have a finite horizon. On the contrary, *step-based* approaches [e.g., 42, 44, 23], derived from the Policy Gradient Theorem [51], can estimate the gradient by averaging over timesteps. The latter requires a function approximator (a critic) to estimate the Q-function, or directly the advantage function [43]. When coming to the on/off-policy dichotomy, the previous distinction has a relevant impact. Indeed, episode-based approaches need to perform importance sampling on trajectories, thus the importance weights are the products of policy ratios for all executed actions within a trajectory, whereas step-based algorithms need just to weight each sample with the corresponding policy ratio. The latter case helps to keep the value of the importance weights close to one, but the need to have a critic prevents from a complete analysis of the uncertainty since the bias/variance injected by the critic is hard to compute [21]. Moreover, in the off-policy scenario, it is necessary to control some notion of dissimilarity between the behavioral and target policy, as the variance increases when moving too far. This is the case of TRPO [42], where the regularization constraint based on the Kullback-Leibler divergence helps controlling the importance weights but originates from an exact bound on the performance improvement. Intuitively, the same rationale applies to the truncation of the importance weights, employed by PPO, that avoids performing too large steps in the policy space. Nevertheless, the step size in TRPO and the truncation range ϵ in PPO are just hyperparameters and have a limited statistical meaning. On the contrary, other actor-critic architectures have been proposed including also experience replay methods, like [61] in which the importance weights are truncated, but the method is able to account for the injected bias. The authors propose to keep a running mean of the best policies seen so far to avoid a hard constraint on the policy dissimilarity. Differently from these methods, POIS directly models the uncertainty due to the importance sampling procedure. The bound in Theorem 4.1 introduces the unique hyperparameter δ which has a precise statistical meaning as confidence level. The optimal value of δ (like the step size in TRPO and ϵ in PPO) is task-dependent and might vary during the learning procedure. Furthermore, POIS is an episode-based approach in which the importance weights account for the whole trajectory at once; this might prevent from assigning credit to valuable subtrajectories (like in the case of Inverted Pendulum, see Figure 1). A possible solution is to resort to per-decision importance sampling [36].

Table 2: Comparison of some policy optimization algorithms according to different dimensions. For brevity, we will indicate with $w_{\theta'}/\theta(a_t|s_t) = \frac{\pi_{\theta'}(a_t|s_t)}{\pi_{\theta}(a_t|s_t)}$. For episode-based algorithms we will indicate with $\hat{\mathbb{E}}_{\tau \sim \theta}$ the empirical average over trajectories collected with π_{θ} . For step-based algorithms $\hat{\mathbb{E}}_{t \sim \theta}$ is the empirical average collecting samples with π_{θ} . For parameter-based algorithms we indicate with $\hat{\mathbb{E}}_{\theta \sim \rho, \tau \sim \theta}$ the empirical expectation taken w.r.t. policy parameter θ sampled from the hyperpolicy ν_{ρ} and trajectory τ collected with π_{θ} . For the actor-critic architectures, \hat{Q} and \hat{A} are the estimated Q-function and advantage function.

Algorithm	Action/Parameter based	On/Off policy	Optimization problem	Critic	Timestep/Trajectory based
REINFORCE/ G(PO)MDP [64, 4]	action-based	on-policy	$\max_{\hat{\mathbb{E}}_{\tau \sim \theta}} [R(\tau)]$	No	episode-based
TRPO [42]	action-based	on-policy	$\max_{\hat{\mathbb{E}}_{t \sim \theta}} [w_{\theta'}/\theta(a_t s_t)\hat{A}(s_t, a_t)]$ s.t. $\hat{\mathbb{E}}_{t \sim \theta} [D_{\text{KL}}(\pi_{\theta'}(\cdot s_t) \parallel \pi_{\theta}(\cdot s_t))] \leq \delta$	Yes	step-based
PPO [44]	action-based	on/off-policy	$\max_{\hat{\mathbb{E}}_{t \sim \theta}} [\min \{w_{\theta'}/\theta(a_t s_t)\hat{A}(s_t, a_t), \text{clip}(w_{\theta'}/\theta(a_t s_t), 1 - \epsilon, 1 + \epsilon)\hat{A}(s_t, a_t)\}]$	Yes	step-based
DDPG [23]	action-based	off-policy	$\max_{\hat{\mathbb{E}}_{t \sim \theta}} [\pi_{\theta'}(a_t s_t)\hat{Q}(s_t, a_t)]$	Yes	step-based
REPS [31] ⁸	action-based	on-policy	$\max_{\hat{\mathbb{E}}_{t \sim \theta}} [R(s_t, a_t)]$ s.t. $\hat{\mathbb{E}}_{t \sim \theta} [D_{\text{KL}}(d_{\mu}^{\pi_{\theta'}}(s_t, a_t) \parallel d_{\mu}^{\pi_{\theta}}(s_t, a_t))] \leq \delta$	Yes	step-based
RWR [32]	action-based	on-policy	$\max_{\hat{\mathbb{E}}_{t \sim \theta}} [\beta \exp(-\beta R(s_t, a_t))]$	No	step-based
A-POIS	action-based	on/off-policy	$\max_{\hat{\mathbb{E}}_{\tau \sim \theta}} [w_{\theta'}/\theta(\tau)R(\tau)] - \lambda \sqrt{\hat{d}_2(p(\cdot \theta) \parallel p(\cdot \theta))}/N$	No	episode-based
PGPE [45]	parameter-based	on-policy	$\max_{\hat{\mathbb{E}}_{\theta \sim \rho, \tau \sim \theta}} [R(\tau)]$	No	episode-based
IW-PGPE [67]	parameter-based	on/off-policy	$\max_{\hat{\mathbb{E}}_{\theta \sim \rho, \tau \sim \theta}} [w_{\rho'}/\rho(\theta)R(\tau)]$	No	episode-based
P-POIS	parameter-based	on/off-policy	$\max_{\hat{\mathbb{E}}_{\theta \sim \rho, \tau \sim \theta}} [w_{\rho'}/\rho(\theta)R(\tau)] - \lambda \sqrt{d_2(\nu_{\rho'} \parallel \nu_{\rho})}/N$	No	episode-based

⁸We indicate with $d_{\mu}^{\pi_{\theta}}(s_t, a_t)$ the state-action occupancy [51].

B Proofs and Derivations

Lemma 4.1. *Let P and Q be two probability measures on the measurable space $(\mathcal{X}, \mathcal{F})$ such that $P \ll Q$. Let $\mathbf{x} = (x_1, x_2, \dots, x_N)^T$ i.i.d. random variables sampled from Q and $f : \mathcal{X} \rightarrow \mathbb{R}$ be a bounded function ($\|f\|_\infty < +\infty$). Then, for any $N > 0$, the variance of the IS estimator $\widehat{\mu}_{P/Q}$ can be upper bounded as:*

$$\mathbb{V}\text{ar}_{\mathbf{x} \sim Q} [\widehat{\mu}_{P/Q}] \leq \frac{1}{N} \|f\|_\infty^2 d_2(P\|Q). \quad (7)$$

Proof. From the fact that x_i are i.i.d. we can write:

$$\begin{aligned} \mathbb{V}\text{ar}_{\mathbf{x} \sim Q} [\widehat{\mu}_{P/Q}] &\leq \frac{1}{N} \mathbb{V}\text{ar}_{x_1 \sim Q} \left[\frac{p(x_1)}{q(x_1)} f(x_1) \right] \leq \frac{1}{N} \mathbb{E}_{x_1 \sim Q} \left[\left(\frac{p(x_1)}{q(x_1)} f(x_1) \right)^2 \right] \\ &\leq \frac{1}{N} \|f\|_\infty^2 \mathbb{E}_{x_1 \sim Q} \left[\left(\frac{p(x_1)}{q(x_1)} \right)^2 \right] = \frac{1}{N} \|f\|_\infty^2 d_2(P\|Q). \end{aligned}$$

□

Theorem 4.1. *Let P and Q be two probability measures on the measurable space $(\mathcal{X}, \mathcal{F})$ such that $P \ll Q$ and $d_2(P\|Q) < +\infty$. Let x_1, x_2, \dots, x_N be i.i.d. random variables sampled from Q , and $f : \mathcal{X} \rightarrow \mathbb{R}$ be a bounded function ($\|f\|_\infty < +\infty$). Then, for any $0 < \delta \leq 1$ and $N > 0$ with probability at least $1 - \delta$ it holds that:*

$$\mathbb{E}_{x \sim P} [f(x)] \geq \frac{1}{N} \sum_{i=1}^N w_{P/Q}(x_i) f(x_i) - \|f\|_\infty \sqrt{\frac{(1-\delta)d_2(P\|Q)}{\delta N}}. \quad (8)$$

Proof. We start from Cantelli's inequality applied on the random variable $\widehat{\mu}_{P/Q} = \frac{1}{N} \sum_{i=1}^N w_{P/Q}(x_i) f(x_i)$:

$$\Pr \left(\widehat{\mu}_{P/Q} - \mathbb{E}_{x \sim P} [f(x)] \geq \lambda \right) \leq \frac{1}{1 + \frac{\lambda^2}{\mathbb{V}\text{ar}_{\mathbf{x} \sim Q} [\widehat{\mu}_{P/Q}]}}. \quad (11)$$

By calling $\delta = \frac{1}{1 + \frac{\lambda^2}{\mathbb{V}\text{ar}_{\mathbf{x} \sim Q} [\widehat{\mu}_{P/Q}]}}$ and considering the complementary event, we get that with probability at least $1 - \delta$ we have:

$$\mathbb{E}_{x \sim P} [f(x)] \geq \widehat{\mu}_{P/Q} - \sqrt{\frac{1-\delta}{\delta} \mathbb{V}\text{ar}_{\mathbf{x} \sim Q} [\widehat{\mu}_{P/Q}]}. \quad (12)$$

By replacing the variance with the bound in Theorem 4.1 we get the result. □

Theorem 4.2. *Let p_ω be a p.d.f. differentiable w.r.t. $\omega \in \Omega$. Then, it holds that, for the Rényi divergence: $D_\alpha(p_{\omega'} \| p_\omega) = \frac{\alpha}{2} (\omega' - \omega)^T \mathcal{F}(\omega) (\omega' - \omega) + o(\|\omega' - \omega\|_2^2)$, and for the exponentiated Rényi divergence: $d_\alpha(p_{\omega'} \| p_\omega) = 1 + \frac{\alpha}{2} (\omega' - \omega)^T \mathcal{F}(\omega) (\omega' - \omega) + o(\|\omega' - \omega\|_2^2)$.*

Proof. We need to compute the second-order Taylor expansion of the α -Rényi divergence. We start considering the term:

$$I(\omega') = \int_{\mathcal{X}} \left(\frac{p_{\omega'}(x)}{p_\omega(x)} \right)^\alpha p_\omega(x) dx = \int_{\mathcal{X}} p_{\omega'}(x)^\alpha p_\omega(x)^{1-\alpha} dx. \quad (13)$$

The gradient is given by:

$$\nabla_{\omega'} I(\omega') = \int_{\mathcal{X}} \nabla_{\omega'} p_{\omega'}(x)^\alpha p_\omega(x)^{1-\alpha} dx = \alpha \int_{\mathcal{X}} p_{\omega'}(x)^{\alpha-1} p_\omega(x)^{1-\alpha} \nabla_{\omega'} p_{\omega'}(x) dx.$$

Thus, $\nabla_{\omega'} I(\omega')|_{\omega'=\omega} = \mathbf{0}$. We now compute the Hessian:

$$\begin{aligned} \mathcal{H}_{\omega'} I(\omega') &= \nabla_{\omega'} \nabla_{\omega'}^T I(\omega') = \alpha \nabla_{\omega'} \int_{\mathcal{X}} p_{\omega'}(x)^{\alpha-1} p_\omega(x)^{1-\alpha} \nabla_{\omega'}^T p_{\omega'}(x) dx \\ &= \alpha \int_{\mathcal{X}} \left((\alpha-1) p_{\omega'}(x)^{\alpha-2} p_\omega(x)^{1-\alpha} \nabla_{\omega'} p_{\omega'}(x) \nabla_{\omega'}^T p_{\omega'}(x) + p_{\omega'}(x)^{\alpha-1} p_\omega(x)^{1-\alpha} \mathcal{H}_{\omega'} p_{\omega'}(x) \right) dx. \end{aligned}$$

Evaluating the Hessian in ω we have:

$$\mathcal{H}_{\omega'} I(\omega')|_{\omega'=\omega} = \alpha(\alpha-1) \int_{\mathcal{X}} p_\omega(x)^{-1} \nabla_\omega p_\omega(x) \nabla_\omega^T p_\omega(x) dx$$

$$= \alpha(\alpha - 1) \int_{\mathcal{X}} p_{\omega}(x) \nabla_{\omega} \log p_{\omega}(x) \nabla_{\omega}^T \log p_{\omega}(x) dx = \alpha(\alpha - 1) \mathcal{F}(\omega).$$

Now, $D_{\alpha}(p_{\omega'} \| p_{\omega}) = \frac{1}{\alpha-1} \log I(\omega')$. Thus:

$$\begin{aligned} \nabla_{\omega'} D_{\alpha}(p_{\omega'} \| p_{\omega})|_{\omega'=\omega} &= \frac{1}{\alpha-1} \frac{\nabla_{\omega'} I(\omega')}{I(\omega')} \Big|_{\omega'=\omega} = \mathbf{0}, \\ \mathcal{H}_{\omega'} D_{\alpha}(p_{\omega'} \| p_{\omega})|_{\omega'=\omega} &= \frac{1}{\alpha-1} \frac{I(\omega') \mathcal{H}_{\omega'} I(\omega') + \nabla_{\omega'} I(\omega') \nabla_{\omega'}^T I(\omega')}{(I(\omega'))^2} \Big|_{\omega'=\omega} \\ &= \frac{1}{\alpha-1} \mathcal{H}_{\omega'} I(\omega')|_{\omega'=\omega} = \alpha \mathcal{F}(\omega), \end{aligned}$$

having observed that $I(\omega) = 1$. For what concerns the $d_{\alpha}(p_{\omega'} \| p_{\omega})$, we have:

$$\begin{aligned} \nabla_{\omega'} d_{\alpha}(p_{\omega'} \| p_{\omega})|_{\omega'=\omega} &= \nabla_{\omega'} \exp(D_{\alpha}(p_{\omega'} \| p_{\omega}))|_{\omega'=\omega} \\ &= \exp(D_{\alpha}(p_{\omega'} \| p_{\omega})) \nabla_{\omega'} D_{\alpha}(p_{\omega'} \| p_{\omega})|_{\omega'=\omega} = \mathbf{0}, \\ \mathcal{H}_{\omega'} d_{\alpha}(p_{\omega'} \| p_{\omega})|_{\omega'=\omega} &= \mathcal{H}_{\omega'} \exp(D_{\alpha}(p_{\omega'} \| p_{\omega}))|_{\omega'=\omega} \\ &= \exp(D_{\alpha}(p_{\omega'} \| p_{\omega})) \left(\mathcal{H}_{\omega'} D_{\alpha}(p_{\omega'} \| p_{\omega}) + \nabla_{\omega'} D_{\alpha}(p_{\omega'} \| p_{\omega}) \nabla_{\omega'}^T D_{\alpha}(p_{\omega'} \| p_{\omega}) \right) \Big|_{\omega'=\omega} \\ &= \alpha \mathcal{F}(\omega). \end{aligned}$$

□

C Analysis of the IS estimator

In this Appendix, we analyze the behavior of the importance weights when the behavioral and target distributions are Gaussians. We start providing a closed-form expression for the Rényi divergence between multivariate Gaussian distributions [6]. Let $P \sim \mathcal{N}(\mu_P, \Sigma_P)$ and $Q \sim \mathcal{N}(\mu_Q, \Sigma_Q)$ and $\alpha \in [0, \infty]$:

$$D_{\alpha}(P \| Q) = \frac{\alpha}{2} (\mu_P - \mu_Q)^T \Sigma_{\alpha}^{-1} (\mu_P - \mu_Q) - \frac{1}{2(\alpha-1)} \log \frac{\det(\Sigma_{\alpha})}{\det(\Sigma_P)^{1-\alpha} \det(\Sigma_Q)^{\alpha}}, \quad (14)$$

where $\Sigma_{\alpha} = \alpha \Sigma_Q + (1-\alpha) \Sigma_P$ under the assumption that Σ_{α} is positive-definite.

From now on, we will focus on univariate Gaussian distributions and we provide a closed-form expression for the importance weights and their probability density function f_w . We consider $Q \sim \mathcal{N}(\mu_Q, \sigma_Q^2)$ as behavioral distribution and $P \sim \mathcal{N}(\mu_P, \sigma_P^2)$ as target distribution. We assume that $\sigma_Q^2, \sigma_P^2 > 0$ and we consider the two cases: unequal variances and equal variances. For brevity, we will indicate with $w(x)$ the weight $w_{P/Q}(x)$.

C.1 Unequal variances

When $\sigma_Q^2 \neq \sigma_P^2$, the expression of the importance weights is given by:

$$w(x) = \frac{\sigma_Q}{\sigma_P} \exp\left(\frac{1}{2} \frac{(\mu_P - \mu_Q)^2}{\sigma_Q^2 - \sigma_P^2}\right) \exp\left(-\frac{1}{2} \frac{\sigma_Q^2 - \sigma_P^2}{\sigma_Q^2 \sigma_P^2} \left(x - \frac{\sigma_Q^2 \mu_P - \sigma_P^2 \mu_Q}{\sigma_Q^2 - \sigma_P^2}\right)^2\right), \quad (15)$$

for $x \sim Q$. Let us first notice two distinct situations: if $\sigma_Q^2 - \sigma_P^2 > 0$ the weight $w(x)$ is upper bounded by $A = \frac{\sigma_Q}{\sigma_P} \exp\left(\frac{1}{2} \frac{(\mu_P - \mu_Q)^2}{\sigma_Q^2 - \sigma_P^2}\right)$, whereas if $\sigma_Q^2 - \sigma_P^2 < 0$, $w(x)$ is unbounded but it admits a minimum of value A . Let us investigate the probability density function.

Proposition C.1. *Let $Q \sim \mathcal{N}(\mu_Q, \sigma_Q^2)$ be the behavioral distribution and $P \sim \mathcal{N}(\mu_P, \sigma_P^2)$ be the target distribution, with $\sigma_Q^2 \neq \sigma_P^2$. The probability density function of $w(x) = p(x)/q(x)$ is given by:*

$$f_w(y) = \begin{cases} \frac{\bar{\sigma}}{y \sqrt{\pi \log \frac{A}{y}}} \exp\left(-\frac{1}{2} \bar{\mu}^2\right) \left(\frac{y}{A}\right)^{\bar{\sigma}^2} \cosh\left(\bar{\mu} \bar{\sigma} \sqrt{2 \log \frac{A}{y}}\right), & \text{if } \sigma_Q^2 > \sigma_P^2, y \in [0, A], \\ \frac{\bar{\sigma}}{y \sqrt{\pi \log \frac{y}{A}}} \exp\left(-\frac{1}{2} \bar{\mu}^2\right) \left(\frac{A}{y}\right)^{\bar{\sigma}^2} \cosh\left(\bar{\mu} \bar{\sigma} \sqrt{2 \log \frac{y}{A}}\right), & \text{if } \sigma_Q^2 < \sigma_P^2, y \in [A, \infty), \end{cases}$$

where $\bar{\mu} = \frac{\sigma_Q}{\sigma_Q^2 - \sigma_P^2} (\mu_P - \mu_Q)$ and $\bar{\sigma}^2 = \frac{\sigma_P^2}{|\sigma_Q^2 - \sigma_P^2|}$.

Proof. We look at $w(x)$ as a function of random variable $x \sim Q$. We introduce the following symbols:

$$m = \frac{\sigma_Q^2 \mu_P - \sigma_P^2 \mu_Q}{\sigma_Q^2 - \sigma_P^2}, \quad \tau = \frac{\sigma_Q^2 - \sigma_P^2}{\sigma_Q^2 \sigma_P^2}.$$

Let us start computing the c.d.f.:

$$F_w(y) = \Pr(w(x) \leq y) = \Pr\left(A \exp\left(-\frac{1}{2}\tau(x-m)^2\right) \leq y\right) = \Pr\left(\tau(x-m)^2 \geq -2 \log \frac{y}{A}\right).$$

We distinguish the two cases according to the sign of τ and we observe that $x = \mu_Q + \sigma_Q z$ where $z \sim \mathcal{N}(0, 1)$:

$\tau > 0$:

$$\begin{aligned} F_w(y) &= \Pr\left((x-m)^2 \geq \frac{2}{\tau} \log \frac{A}{y}\right) \\ &= \Pr\left(x \leq m - \sqrt{\frac{2}{\tau} \log \frac{A}{y}}\right) + \Pr\left(x \geq m + \sqrt{\frac{2}{\tau} \log \frac{A}{y}}\right) \\ &= \Pr\left(z \leq \frac{m - \mu_Q}{\sigma_Q} - \sqrt{\frac{2}{\tau \sigma_Q^2} \log \frac{A}{y}}\right) + \Pr\left(z \geq \frac{m - \mu_Q}{\sigma_Q} + \sqrt{\frac{2}{\tau \sigma_Q^2} \log \frac{A}{y}}\right). \end{aligned}$$

We call $\bar{\mu} = \frac{m - \mu_Q}{\sigma_Q} = \frac{\sigma_Q}{\sigma_Q^2 - \sigma_P^2}(\mu_P - \mu_Q)$ and $\bar{\sigma}^2 = \frac{1}{\tau \sigma_Q^2} = \frac{\sigma_P^2}{\sigma_Q^2 - \sigma_P^2}$, thus we have:

$$\begin{aligned} F_w(y) &= \Pr\left(z \leq \bar{\mu} - \sqrt{2\bar{\sigma}^2 \log \frac{A}{y}}\right) + \Pr\left(z \geq \bar{\mu} + \sqrt{2\bar{\sigma}^2 \log \frac{A}{y}}\right) \\ &= \Phi\left(\bar{\mu} - \sqrt{2\bar{\sigma}^2 \log \frac{A}{y}}\right) + 1 - \Phi\left(\bar{\mu} + \sqrt{2\bar{\sigma}^2 \log \frac{A}{y}}\right), \end{aligned}$$

where Φ is the c.d.f. of a normal standard distribution. By taking the derivative w.r.t. y we get the p.d.f.:

$$\begin{aligned} f_w(y) &= \frac{\partial F_w(y)}{\partial y} = -\sqrt{2\bar{\sigma}^2} \frac{1}{2\sqrt{\log \frac{A}{y}}} \frac{y}{A} \frac{-A}{y^2} \left(\phi\left(\bar{\mu} - \sqrt{2\bar{\sigma}^2 \log \frac{A}{y}}\right) + \phi\left(\bar{\mu} + \sqrt{2\bar{\sigma}^2 \log \frac{A}{y}}\right)\right) \\ &= \frac{\sqrt{2\bar{\sigma}}}{2y\sqrt{\log \frac{A}{y}}} \left(\phi\left(\bar{\mu} - \sqrt{2\bar{\sigma}^2 \log \frac{A}{y}}\right) + \phi\left(\bar{\mu} + \sqrt{2\bar{\sigma}^2 \log \frac{A}{y}}\right)\right) \\ &= \frac{\sqrt{2\bar{\sigma}}}{2y\sqrt{\log \frac{A}{y}}} \frac{1}{\sqrt{2\pi}} \left(\exp\left(-\frac{1}{2}\left(\bar{\mu} - \sqrt{2\bar{\sigma}^2 \log \frac{A}{y}}\right)^2\right) + \exp\left(-\frac{1}{2}\left(\bar{\mu} + \sqrt{2\bar{\sigma}^2 \log \frac{A}{y}}\right)^2\right)\right) \\ &= \frac{\bar{\sigma}}{y\sqrt{\pi \log \frac{A}{y}}} \exp\left(-\frac{1}{2}\bar{\mu}^2\right) \exp\left(-\bar{\sigma}^2 \log \frac{A}{y}\right) \frac{\exp\left(\bar{\mu}\bar{\sigma}\sqrt{2 \log \frac{A}{y}}\right) + \exp\left(-\bar{\mu}\bar{\sigma}\sqrt{2 \log \frac{A}{y}}\right)}{2} \\ &= \frac{\bar{\sigma}}{y\sqrt{\pi \log \frac{A}{y}}} \exp\left(-\frac{1}{2}\bar{\mu}^2\right) \left(\frac{y}{A}\right)^{\bar{\sigma}^2} \cosh\left(\bar{\mu}\bar{\sigma}\sqrt{2 \log \frac{A}{y}}\right), \end{aligned}$$

where ϕ is the p.d.f. of a normal standard distribution.

$\tau < 0$: The derivation takes similar steps, all it takes is to call $\bar{\sigma}^2 = -\frac{1}{\tau \sigma_Q^2} = \frac{\sigma_P^2}{\sigma_P^2 - \sigma_Q^2}$, then the c.d.f. becomes:

$$F_w(y) = \Phi\left(\bar{\mu} + \sqrt{2\bar{\sigma}^2 \log \frac{y}{A}}\right) - \Phi\left(\bar{\mu} - \sqrt{2\bar{\sigma}^2 \log \frac{y}{A}}\right),$$

and the p.d.f. is:

$$f_w(x) = \frac{\bar{\sigma}}{y\sqrt{\pi \log \frac{y}{A}}} \exp\left(-\frac{1}{2}\bar{\mu}^2\right) \left(\frac{A}{y}\right)^{\bar{\sigma}^2} \cosh\left(\bar{\mu}\bar{\sigma}\sqrt{2 \log \frac{y}{A}}\right).$$

To unify the two cases we set $\bar{\sigma}^2 = \frac{\sigma_P^2}{|\sigma_Q^2 - \sigma_P^2|}$. □

It is interesting to investigate the properties of the tail of the distribution when w is unbounded. Indeed, we discover that the distribution displays a fat-tail behavior.

Proposition C.2. *If $\sigma_P^2 > \sigma_Q^2$ then there exists $c > 0$ and $y_0 > 0$ such that for any $y \geq y_0$, the p.d.f. f_w can be lower bounded as $f_w(y) \geq cy^{-1-\bar{\sigma}^2}(\log y)^{-\frac{1}{2}}$.*

Proof. Let us call $z = y/A$ and let $a > 0$ be a constant, then it holds that for sufficiently large y we have:

$$f_w(y) \geq az^{-1-\bar{\sigma}^2}(\log z)^{-1/2} \exp\left(\sqrt{\log z}\right)^{\sqrt{2\mu\bar{\sigma}}}. \quad (16)$$

To get the result, we observe that for $z > 1$ we have $\exp\left(\sqrt{\log z}\right) \geq 1$. Now, by replacing z with y/A we just need to change the constant a into $c > 0$. \square

As a consequence, the α -th moment of $w(x)$ does not exist for $\alpha - 1 - \bar{\sigma}^2 \geq -1 \implies \alpha \geq \bar{\sigma}^2 = \frac{\sigma_P^2}{\sigma_P^2 - \sigma_Q^2}$, this prevents from using Bernstein-like inequalities for bounding in probability the importance weights. The non-existence of finite moments is confirmed by the α -Rényi divergence. Indeed, the α -Rényi divergence is defined when $\sigma_\alpha^2 = \alpha\sigma_Q^2 + (1 - \alpha)\sigma_P^2 > 0$, i.e., $\alpha < \frac{\sigma_P^2}{\sigma_P^2 - \sigma_Q^2}$.

C.2 Equal variances

If $\sigma_Q^2 = \sigma_P^2 = \sigma^2$, the importance weights have the following expression:

$$w(x) = \exp\left(\frac{\mu_P - \mu_Q}{\sigma^2} \left(x - \frac{\mu_P + \mu_Q}{2}\right)\right), \quad (17)$$

for $x \sim Q$. The weight $w(x)$ is clearly unbounded and has 0 as infimum value. Let us investigate its probability density function.

Proposition C.3. *Let $Q \sim \mathcal{N}(\mu_Q, \sigma^2)$ be the behavioral distribution and $P \sim \mathcal{N}(\mu_P, \sigma^2)$ be the target distribution. The probability density function of $w(x) = q(x)/p(x)$ is given by:*

$$f_w(y) = \frac{|\tilde{\sigma}|}{\sqrt{2\pi}y^{\frac{3}{2}}} \exp\left(-\frac{1}{2}(\tilde{\mu}^2 + \tilde{\sigma}^2(\log y)^2)\right), \quad (18)$$

where $\tilde{\mu} = \frac{\mu_P - \mu_Q}{2\sigma}$ and $\tilde{\sigma} = \frac{\sigma}{\mu_P - \mu_Q}$.

Proof. We start computing the c.d.f.:

$$F_w(y) = \Pr\left(\exp\left\{\frac{\mu_P - \mu_Q}{\sigma^2} \left(x - \frac{\mu_P + \mu_Q}{2}\right)\right\} \leq y\right) = \Pr\left(\frac{\mu_P - \mu_Q}{\sigma^2} \left(x - \frac{\mu_P + \mu_Q}{2}\right) \leq \log y\right).$$

First, we consider the case $\mu_P - \mu_Q > 0$ and observe that $x = \mu_Q + \sigma z$, where $z \sim \mathcal{N}(0, 1)$:

$$F_w(y) = \Pr\left(x \leq \frac{\mu_P + \mu_Q}{2} + \frac{\sigma^2}{\mu_P - \mu_Q} \log y\right) = \Pr\left(z \leq \frac{\mu_P - \mu_Q}{2\sigma} + \frac{\sigma}{\mu_P - \mu_Q} \log y\right).$$

We call $\tilde{\mu} = \frac{\mu_P - \mu_Q}{2\sigma}$ and $\tilde{\sigma} = \frac{\sigma}{\mu_P - \mu_Q}$ and we have:

$$F_w(y) = \Pr(z \leq \tilde{\mu} + \tilde{\sigma} \log y) = \Phi(\tilde{\mu} + \tilde{\sigma} \log y).$$

We take the derivative in order to get the density function:

$$f_w(y) = \frac{\partial F_w(y)}{\partial y} = \frac{\tilde{\sigma}}{y} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}(\tilde{\mu} + \tilde{\sigma} \log y)^2\right) = \frac{\tilde{\sigma}}{\sqrt{2\pi}y^{\tilde{\mu}\tilde{\sigma}+1}} \exp\left(-\frac{1}{2}(\tilde{\mu}^2 + \tilde{\sigma}^2(\log y)^2)\right).$$

For the case $\mu_P - \mu_Q < 0$ the derivation is symmetric and the p.d.f. differs only by a minus sign. We account for this fact by considering $|\tilde{\sigma}|$ in the final formula. \square

In the case of equal variances, the tail behavior is different.

Proposition C.4. *If $\sigma_P^2 = \sigma_Q^2$ then for any $\alpha > 0$ there exist $c > 0$ and $y_0 > 0$ such that for any $y \geq y_0$, the p.d.f. can be upper bounded as $f_w(y) \leq cy^{-\alpha}$.*

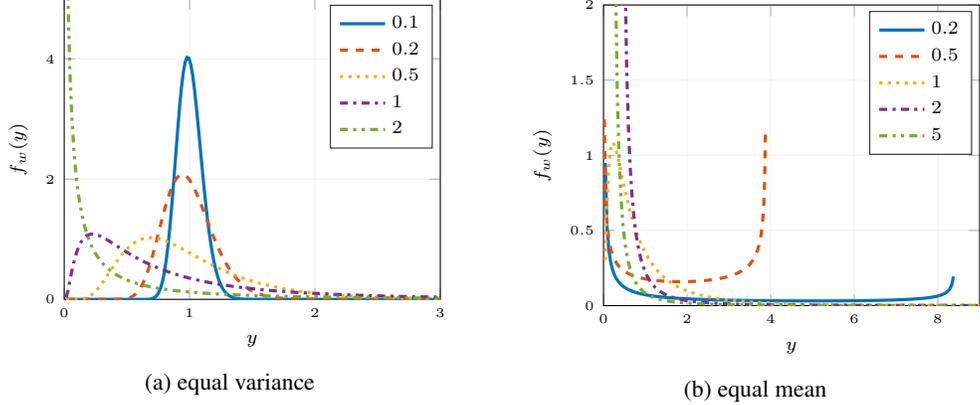


Figure 3: Probability density function of the importance weights when the behavioral distribution is $\mathcal{N}(0, 1)$ and the mean is changed keeping the variance equal to 1 (a) or the variance is changed keeping the target mean equal to 1 (b).

Proof. Condensing all the constants in c , the p.d.f. can be written as:

$$f_w(y) = cy^{-3/2} \exp((\log y)^2)^{-\frac{\alpha}{2}}. \quad (19)$$

For any $\alpha > 0$, let us solve the following inequality:

$$y^{3/2} \exp((\log y)^2)^{\frac{\alpha}{2}} \geq y^\alpha \implies y \geq \exp\left(\frac{2}{\alpha} \left(\alpha - \frac{3}{2}\right)\right). \quad (20)$$

Thus, for $y \geq \exp\left(\frac{2}{\alpha} \left(\alpha - \frac{3}{2}\right)\right)$ we have that $f_w(y) \leq cy^{-\alpha}$. \square

This is sufficient to ensure the existence of the moments of any order, indeed the corresponding Rényi divergence is: $\frac{\alpha(\mu_P - \mu_Q)^2}{2\sigma^2}$. By the way, the distribution of $w(x)$ remains subexponential, as $\exp\left((\log y)^2\right)^{-\frac{\alpha}{2}} \geq e^{-\eta y}$ for sufficiently large y .

Figure 3 reports the p.d.f. of the importance weights for different values of mean and variance of the target distribution.

D Analysis of the SN Estimator

In this Appendix, we provide some results regarding bias and variance of the self-normalized importance sampling estimator. Let us start with the following result, derived from [8], that bounds the expected squared difference between non-self-normalized weight $w(x)$ and self-normalized weight $\tilde{w}(x)$.

Lemma D.1. *Let P and Q be two probability measures on the measurable space $(\mathcal{X}, \mathcal{F})$ such that $P \ll Q$ and $d_2(P\|Q) < +\infty$. Let x_1, x_2, \dots, x_N i.i.d. random variables sampled from Q . Then, for $N > 0$ and for any $i = 1, 2, \dots, N$ it holds that:*

$$\mathbb{E}_{\mathbf{x} \sim Q} \left[\left(\tilde{w}_{P/Q}(x_i) - \frac{w_{P/Q}(x_i)}{N} \right)^2 \right] \leq \frac{d_2(P\|Q) - 1}{N}. \quad (21)$$

Proof. The result derives from simple algebraic manipulations and from the fact that $\text{Var}_{\mathbf{x} \sim Q} [w_{P/Q}(x)] = d_2(P\|Q) - 1$.

$$\begin{aligned} \mathbb{E}_{\mathbf{x} \sim Q} \left[\left(\tilde{w}_{P/Q}(x_i) - \frac{w_{P/Q}(x_i)}{N} \right)^2 \right] &= \mathbb{E}_{\mathbf{x} \sim Q} \left[\left(\frac{w_{P/Q}(x_i)}{\sum_{j=1}^N w_{P/Q}(x_j)} \right)^2 \left(1 - \frac{\sum_{j=1}^N w_{P/Q}(x_j)}{N} \right)^2 \right] \\ &\leq \mathbb{E}_{\mathbf{x} \sim Q} \left[\left(1 - \frac{\sum_{j=1}^N w_{P/Q}(x_j)}{N} \right)^2 \right] = \text{Var}_{\mathbf{x} \sim Q} \left[\frac{\sum_{j=1}^N w_{P/Q}(x_j)}{N} \right] \end{aligned}$$

$$= \frac{1}{N} \mathbb{V}\text{ar}_{x_1 \sim Q} [w_{P/Q}(x_1)] = \frac{d_2(P\|Q) - 1}{N}.$$

□

A similar argument can be used to derive a bound on the bias of the SN estimator.

Proposition D.1. *Let P and Q be two probability measures on the measurable space $(\mathcal{X}, \mathcal{F})$ such that $P \ll Q$ and $d_2(P\|Q) < +\infty$. Let x_1, x_2, \dots, x_N i.i.d. random variables sampled from Q and $f : \mathcal{X} \rightarrow \mathbb{R}$ be a bounded function ($\|f\|_\infty < \infty$). Then, the bias of the SN estimator can be bounded as:*

$$\left| \mathbb{E}_{\mathbf{x} \sim Q} [\tilde{\mu}_{P/Q}] - \mathbb{E}_{x \sim P} [f(x)] \right| \leq \|f\|_\infty \min \left\{ 2, \sqrt{\frac{d_2(P\|Q) - 1}{N}} \right\}. \quad (22)$$

Proof. Since it holds that $|\tilde{\mu}_{P/Q}| \leq \|f\|_\infty$ the bias cannot be larger than $2\|f\|_\infty$. We now derive a bound for the bias that vanishes as $N \rightarrow \infty$. We exploit the fact that the IS estimator is unbiased, i.e., $\mathbb{E}_{\mathbf{x} \sim Q} [\hat{\mu}_{P/Q}] = \mathbb{E}_{x \sim P} [f(x)]$.

$$\begin{aligned} \left| \mathbb{E}_{\mathbf{x} \sim Q} [\tilde{\mu}_{P/Q}] - \mathbb{E}_{x \sim P} [f(x)] \right| &= \left| \mathbb{E}_{\mathbf{x} \sim Q} [\tilde{\mu}_{P/Q}] - \mathbb{E}_{\mathbf{x} \sim Q} [\hat{\mu}_{P/Q}] \right| = \left| \mathbb{E}_{\mathbf{x} \sim Q} [\tilde{\mu}_{P/Q} - \hat{\mu}_{P/Q}] \right| \\ &\leq \mathbb{E}_{\mathbf{x} \sim Q} [|\tilde{\mu}_{P/Q} - \hat{\mu}_{P/Q}|] = \\ &= \mathbb{E}_{\mathbf{x} \sim Q} \left[\left| \frac{\sum_{i=1}^N w_{P/Q}(x_i) f(x_i)}{\sum_{i=1}^N w_{P/Q}(x_i)} - \frac{\sum_{i=1}^N w_{P/Q}(x_i) f(x_i)}{N} \right| \right] \\ &= \mathbb{E}_{\mathbf{x} \sim Q} \left[\left| \frac{\sum_{i=1}^N w_{P/Q}(x_i) f(x_i)}{\sum_{i=1}^N w_{P/Q}(x_i)} \right| \left| 1 - \frac{\sum_{i=1}^N w_{P/Q}(x_i)}{N} \right| \right] \end{aligned} \quad (23)$$

$$\leq \mathbb{E}_{\mathbf{x} \sim Q} \left[\left(\frac{\sum_{i=1}^N w_{P/Q}(x_i) f(x_i)}{\sum_{i=1}^N w_{P/Q}(x_i)} \right)^2 \right]^{\frac{1}{2}} \mathbb{E}_{\mathbf{x} \sim Q} \left[\left(1 - \frac{\sum_{i=1}^N w_{P/Q}(x_i)}{N} \right)^2 \right]^{\frac{1}{2}} \quad (24)$$

$$\leq \|f\|_\infty \sqrt{\frac{d_2(P\|Q) - 1}{N}}, \quad (25)$$

where (24) follows from (23) by applying Cauchy-Schwartz inequality and (25) is obtained by observing that $\left(\frac{\sum_{i=1}^N w_{P/Q}(x_i) f(x_i)}{\sum_{i=1}^N w_{P/Q}(x_i)} \right)^2 \leq \|f\|_\infty^2$. □

Bounding the variance of the SN estimator is non-trivial since the the normalization term makes all the samples interdependent. Exploiting the boundedness of $\tilde{\mu}_{P/Q}$ we can derive trivial bounds like: $\mathbb{V}\text{ar}_{\mathbf{x} \sim Q} [\tilde{\mu}_{P/Q}] \leq \|f\|_\infty^2$. However, this bound does not shrink with the number of samples N . Several approximations of the variance have been proposed, like the following derived using the delta method [60, 29]:

$$\mathbb{V}\text{ar}_{\mathbf{x} \sim Q} [\tilde{\mu}_{P/Q}] = \frac{1}{N} \mathbb{E}_{x_1 \sim Q} \left[w_{P/Q}^2(x_1) \left(f(x_1) - \mathbb{E}_{x \sim P} [f(x)] \right)^2 \right] + o(N^{-2}). \quad (26)$$

We will not use the approximate expression for the variance, but we will directly bound the Mean Squared Error (MSE) of the SN estimator, which is the sum of the variance and the bias squared.

Proposition D.2. *Let P and Q be two probability measures on the measurable space $(\mathcal{X}, \mathcal{F})$ such that $P \ll Q$ and $d_2(P\|Q) < +\infty$. Let x_1, x_2, \dots, x_N i.i.d. random variables sampled from Q and $f : \mathcal{X} \rightarrow \mathbb{R}$ be a bounded function ($\|f\|_\infty < +\infty$). Then, the MSE of the SN estimator can be bounded as:*

$$\text{MSE}_{\mathbf{x} \sim Q} [\tilde{\mu}_{P/Q}] \leq 2\|f\|_\infty^2 \min \left\{ 2, \frac{2d_2(P\|Q) - 1}{N} \right\}. \quad (27)$$

Proof. First, recall that $\tilde{\mu}_{P/Q}$ is bounded by $\|f\|_\infty$ thus its MSE cannot be larger than $4\|f\|_\infty^2$. The idea of the proof is to sum and subtract the IS estimator $\hat{\mu}_{P/Q}$:

$$\text{MSE}_{\mathbf{x} \sim Q} [\tilde{\mu}_{P/Q}] = \mathbb{E}_{\mathbf{x} \sim Q} \left[\left(\tilde{\mu}_{P/Q} - \mathbb{E}_{x \sim P} [f(x)] \right)^2 \right]$$

$$= \mathbb{E}_{\mathbf{x} \sim Q} \left[\left(\tilde{\mu}_{P/Q} - \mathbb{E}_{x \sim P} [f(x)] \pm \hat{\mu}_{P/Q} \right)^2 \right] \quad (28)$$

$$\leq 2 \mathbb{E}_{\mathbf{x} \sim Q} \left[\left(\tilde{\mu}_{P/Q} - \hat{\mu}_{P/Q} \right)^2 \right] + 2 \mathbb{E}_{\mathbf{x} \sim Q} \left[\left(\hat{\mu}_{P/Q} - \mathbb{E}_{x \sim P} [f(x)] \right)^2 \right] \quad (29)$$

$$\leq 2 \mathbb{E}_{\mathbf{x} \sim Q} \left[\left(\frac{\sum_{i=1}^N w_{P/Q}(x_i) f(x_i)}{\sum_{i=1}^N w_{P/Q}(x_i)} \right)^2 \left(1 - \frac{\sum_{i=1}^N w_{P/Q}(x_i)}{N} \right)^2 \right] + 2 \mathbb{V}\text{ar}_{\mathbf{x} \sim Q} [\hat{\mu}_{P/Q}] \quad (30)$$

$$\leq 2 \|f\|_\infty^2 \mathbb{E}_{\mathbf{x} \sim Q} \left[\left(1 - \frac{\sum_{i=1}^N w_{P/Q}(x_i)}{N} \right)^2 \right] + 2 \mathbb{V}\text{ar}_{\mathbf{x} \sim Q} [\hat{\mu}_{P/Q}] \quad (31)$$

$$\leq 2 \|f\|_\infty^2 \mathbb{V}\text{ar}_{\mathbf{x} \sim Q} \left[\frac{\sum_{i=1}^N w_{P/Q}(x_i)}{N} \right] + 2 \mathbb{V}\text{ar}_{\mathbf{x} \sim Q} [\hat{\mu}_{P/Q}] \quad (32)$$

$$\leq 2 \|f\|_\infty^2 \frac{d_2(P\|Q) - 1}{N} + 2 \|f\|_\infty^2 \frac{d_2(P\|Q)}{N} = 2 \|f\|_\infty^2 \frac{2d_2(P\|Q) - 1}{N}, \quad (33)$$

where line (29) follows from line (28) by applying the inequality $(a + b)^2 \leq 2(a^2 + b^2)$, (31) follows from (30) by observing that $\left(\frac{\sum_{i=1}^N w_{P/Q}(x_i) f(x_i)}{\sum_{i=1}^N w_{P/Q}(x_i)} \right)^2 \leq \|f\|_\infty^2$. \square

We can use this result to provide a high confidence bound for the SN estimator.

Proposition D.3. *Let P and Q be two probability measures on the measurable space $(\mathcal{X}, \mathcal{F})$ such that $P \ll Q$ and $d_2(P\|Q) < +\infty$. Let x_1, x_2, \dots, x_N i.i.d. random variables sampled from Q and $f : \mathcal{X} \rightarrow \mathbb{R}$ be a bounded function ($\|f\|_\infty < +\infty$). Then, for any $0 < \delta \leq 1$ and $N > 0$ with probability at least $1 - \delta$:*

$$\mathbb{E}_{x \sim P} [f(x)] \geq \frac{1}{N} \sum_{i=1}^N \tilde{w}_{P/Q}(x_i) f(x_i) - 2 \|f\|_\infty \min \left\{ 1, \sqrt{\frac{d_2(P\|Q)(4 - 3\delta)}{\delta N}} \right\}.$$

Proof. The result is obtained by applying Cantelli's inequality and accounting for the bias. Consider the random variable $\tilde{\mu}_{P/Q} = \frac{1}{N} \sum_{i=1}^N \tilde{w}_{P/Q}(x_i) f(x_i)$ and let $\tilde{\lambda} = \lambda - |\mathbb{E}_{x \sim P} [f(x)] - \mathbb{E}_{\mathbf{x} \sim P} [\tilde{\mu}_{P/Q}]|$:

$$\begin{aligned} \Pr \left(\tilde{\mu}_{P/Q} - \mathbb{E}_{x \sim P} [f(x)] \geq \lambda \right) &= \Pr \left(\tilde{\mu}_{P/Q} - \mathbb{E}_{\mathbf{x} \sim P} [\tilde{\mu}_{P/Q}] \geq \lambda + \mathbb{E}_{x \sim P} [f(x)] - \mathbb{E}_{\mathbf{x} \sim P} [\tilde{\mu}_{P/Q}] \right) \\ &\leq \Pr \left(\tilde{\mu}_{P/Q} - \mathbb{E}_{\mathbf{x} \sim P} [\tilde{\mu}_{P/Q}] \geq \lambda - \left| \mathbb{E}_{x \sim P} [f(x)] - \mathbb{E}_{\mathbf{x} \sim P} [\tilde{\mu}_{P/Q}] \right| \right) \\ &= \Pr \left(\tilde{\mu}_{P/Q} - \mathbb{E}_{\mathbf{x} \sim P} [\tilde{\mu}_{P/Q}] \geq \tilde{\lambda} \right). \end{aligned}$$

Now we apply Cantelli's inequality:

$$\begin{aligned} \Pr \left(\tilde{\mu}_{P/Q} - \mathbb{E}_{x \sim P} [f(x)] \geq \lambda \right) &\leq \Pr \left(\tilde{\mu}_{P/Q} - \mathbb{E}_{\mathbf{x} \sim P} [\tilde{\mu}_{P/Q}] \geq \tilde{\lambda} \right) \leq \frac{1}{1 + \frac{\tilde{\lambda}^2}{\mathbb{V}\text{ar}_{\mathbf{x} \sim Q} [\tilde{\mu}_{P/Q}]}} \\ &= \frac{1}{1 + \frac{(\lambda - |\mathbb{E}_{x \sim P} [f(x)] - \mathbb{E}_{\mathbf{x} \sim P} [\tilde{\mu}_{P/Q}]|)^2}{\mathbb{V}\text{ar}_{\mathbf{x} \sim Q} [\tilde{\mu}_{P/Q}]}}. \end{aligned} \quad (34)$$

By calling $\delta = \frac{1}{1 + \frac{(\lambda - |\mathbb{E}_{x \sim P} [f(x)] - \mathbb{E}_{\mathbf{x} \sim P} [\tilde{\mu}_{P/Q}]|)^2}{\mathbb{V}\text{ar}_{\mathbf{x} \sim Q} [\tilde{\mu}_{P/Q}]}}$ and considering the complementary event, we get that with probability at least $1 - \delta$ we have:

$$\mathbb{E}_{x \sim P} [f(x)] \geq \tilde{\mu}_{P/Q} - \left| \mathbb{E}_{x \sim P} [f(x)] - \mathbb{E}_{\mathbf{x} \sim P} [\tilde{\mu}_{P/Q}] \right| - \sqrt{\frac{1 - \delta}{\delta} \mathbb{V}\text{ar}_{\mathbf{x} \sim Q} [\tilde{\mu}_{P/Q}]} \quad (35)$$

Then we bound the bias term $|\mathbb{E}_{x \sim P} [f(x)] - \mathbb{E}_{\mathbf{x} \sim P} [\tilde{\mu}_{P/Q}]|$ with equation (22) and the variance term with the MSE in equation (27). With some simple algebraic manipulation we have:

$$\begin{aligned} \mathbb{E}_{x \sim P} [f(x)] &\geq \tilde{\mu}_{P/Q} - \|f\|_\infty \sqrt{\frac{d_2(P\|Q) - 1}{N}} - \|f\|_\infty \sqrt{\frac{1 - \delta}{\delta} \frac{2(2d_2(P\|Q) - 1)}{N}} \\ &\geq \tilde{\mu}_{P/Q} - \|f\|_\infty \sqrt{\frac{d_2(P\|Q)}{N}} - \|f\|_\infty \sqrt{\frac{1 - \delta}{\delta} \frac{4d_2(P\|Q)}{N}} \end{aligned}$$

$$\begin{aligned}
&= \tilde{\mu}_{P/Q} - \|f\|_\infty \sqrt{\frac{d_2(P\|Q)}{N}} \left(1 + 2\sqrt{\frac{1-\delta}{\delta}}\right) \\
&\geq \tilde{\mu}_{P/Q} - 2\|f\|_\infty \sqrt{\frac{d_2(P\|Q)}{N}} \sqrt{1 + \frac{4(1-\delta)}{\delta}} \\
&\geq \tilde{\mu}_{P/Q} - 2\|f\|_\infty \sqrt{\frac{d_2(P\|Q)(4-3\delta)}{\delta N}},
\end{aligned}$$

where the last line follows from the fact that $\sqrt{a} + \sqrt{b} \leq 2\sqrt{a+b}$ for any $a, b \geq 0$. Finally, recalling that the range of the SN estimator is $2\|f\|_\infty$ we get the result. \square

It is worth noting that, apart for the constants, the bound has the same dependence on d_2 as in Theorem 4.1. Thus, by suitably redefining the hyperparameter λ we can optimize the same surrogate objective function for both IS and SN estimators.

E Implementation details

In this Appendix, we provide some aspects about our implementation of POIS.

E.1 Line Search

At each offline iteration k the parameter update is performed in the direction of $\mathcal{G}(\theta_k^j)^{-1} \nabla_{\theta_k^j} \mathcal{L}(\theta_k^j / \theta_0^j)$ with a step size α_k determined in order to maximize the improvement.

For brevity we will remove subscripts and dependence on θ_0^j from the involved quantities. The rationale behind our line search is the following. Suppose that our objective function $\mathcal{L}(\theta)$, restricted to the gradient direction $\mathcal{G}^{-1}(\theta) \nabla_{\theta} \mathcal{L}(\theta)$, represents a concave parabola in the Riemann manifold having $\mathcal{G}(\theta)$ as Riemann metric tensor. Suppose we know a point θ_0 , the Riemann gradient in that point $\mathcal{G}(\theta_0)^{-1} \nabla_{\theta} \mathcal{L}(\theta_0)$ and another point: $\theta_l = \theta_0 + \alpha_l \mathcal{G}(\theta_0)^{-1} \nabla_{\theta} \mathcal{L}(\theta_0)$. For both points we know the value of the loss function: $\mathcal{L}_0 = \mathcal{L}(\theta_0)$ and $\mathcal{L}_l = \mathcal{L}(\theta_l)$ and indicate with $\Delta \mathcal{L}_l = \mathcal{L}_l - \mathcal{L}_0$ the objective function improvement. Having this information we can compute the vertex of that parabola, which is its global maximum. Let us call $l(\alpha) = \mathcal{L}(\theta_0 + \alpha \mathcal{G}^{-1}(\theta_0) \nabla_{\theta} \mathcal{L}(\theta_0)) - \mathcal{L}(\theta_0)$, being a parabola it can be expressed as $l(\alpha) = a\alpha^2 + b\alpha + c$. Clearly, $c = 0$ by definition of $l(\alpha)$; a and b can be determined by enforcing the conditions:

$$\begin{aligned}
b &= \left. \frac{\partial l}{\partial \alpha} \right|_{\alpha=0} = \left. \frac{\partial}{\partial \alpha} \mathcal{L}(\theta_0 + \alpha \mathcal{G}^{-1}(\theta_0) \nabla_{\theta} \mathcal{L}(\theta_0)) - \mathcal{L}(\theta_0) \right|_{\alpha=0} = \\
&= \nabla_{\theta} \mathcal{L}(\theta_0)^T \mathcal{G}^{-1}(\theta_0) \nabla_{\theta} \mathcal{L}(\theta_0) = \\
&= \|\nabla_{\theta} \mathcal{L}(\theta_0)\|_{\mathcal{G}^{-1}(\theta_0)}^2, \\
l(\alpha_l) &= a\alpha_l^2 + b\alpha_l = a\alpha_l^2 + \|\nabla_{\theta} \mathcal{L}(\theta_0)\|_{\mathcal{G}^{-1}(\theta_0)}^2 \alpha_l = \Delta \mathcal{L}_l \implies \\
\implies a &= \frac{\Delta \mathcal{L}_l - \|\nabla_{\theta} \mathcal{L}(\theta_0)\|_{\mathcal{G}^{-1}(\theta_0)}^2 \alpha_l}{\alpha_l^2}.
\end{aligned}$$

Therefore, the parabola has the form:

$$l(\alpha) = \frac{\Delta \mathcal{L}_l - \|\nabla_{\theta} \mathcal{L}(\theta_0)\|_{\mathcal{G}^{-1}(\theta_0)}^2 \alpha_l}{\alpha_l^2} \alpha^2 + \|\nabla_{\theta} \mathcal{L}(\theta_0)\|_{\mathcal{G}^{-1}(\theta_0)}^2 \alpha. \quad (36)$$

Clearly, the parabola is concave only if $\Delta \mathcal{L}_l < \|\nabla_{\theta} \mathcal{L}(\theta_0)\|_{\mathcal{G}^{-1}(\theta_0)}^2 \alpha_l$. The vertex is located at:

$$\alpha_{l+1} = \frac{\|\nabla_{\theta} \mathcal{L}(\theta_0)\|_{\mathcal{G}^{-1}(\theta_0)}^2 \alpha_l^2}{2 \left(\|\nabla_{\theta} \mathcal{L}(\theta_0)\|_{\mathcal{G}^{-1}(\theta_0)}^2 \alpha_l - \Delta \mathcal{L}_l \right)}. \quad (37)$$

To simplify the expression, like in [25] we define $\alpha_l = \epsilon_l / \|\nabla_{\theta} \mathcal{L}(\theta_0)\|_{\mathcal{G}^{-1}(\theta_0)}^2$. Thus, we get:

$$\epsilon_{l+1} = \frac{\epsilon_l^2}{2(\epsilon_l - \Delta \mathcal{L}_l)}. \quad (38)$$

Of course, we need also to manage the case in which the parabola is convex, i.e., $\Delta\mathcal{L}_l \geq \|\nabla_{\theta}\mathcal{L}(\theta_0)\|_{\mathcal{G}^{-1}(\theta_0)}^2\alpha_l$. Since our objective function is not really a parabola we reinterpret the two cases: i) $\Delta\mathcal{L}_l > \|\nabla_{\theta}\mathcal{L}(\theta_0)\|_{\mathcal{G}^{-1}(\theta_0)}^2\alpha_l$, the function is sublinear and in this case we use (38) to determine the new step size $\alpha_{l+1} = \epsilon_{l+1}/\|\nabla_{\theta}\mathcal{L}(\theta_0)\|_{\mathcal{G}^{-1}(\theta_0)}^2$; ii) $\Delta\mathcal{L}_l \geq \|\nabla_{\theta}\mathcal{L}(\theta_0)\|_{\mathcal{G}^{-1}(\theta_0)}^2\alpha_l$, the function is superlinear, in this case we increase the step size multiplying by $\eta > 1$, i.e., $\alpha_{l+1} = \eta\alpha_l$. Finally the update rule becomes:

$$\epsilon_{l+1} = \begin{cases} \eta\epsilon_l & \text{if } \Delta\mathcal{L}_l > \frac{\epsilon_l(2\eta-1)}{2\eta} \\ \frac{\epsilon_l^2}{2(\epsilon_l-\Delta\mathcal{L}_l)} & \text{otherwise} \end{cases}. \quad (39)$$

The procedure is iterated until a maximum number of attempts is reached (say 30) or the objective function improvement is too small (say $1e-4$). The pseudocode of the line search is reported in Algorithm 3.

Algorithm 3 Parabolic Line Search

Input: $\text{tol}_{\Delta\mathcal{L}} = 1e-4$, $M_{\text{ls}} = 30$, \mathcal{L}_0
Output: α^*

$\alpha_0 = 0$

$\epsilon_1 = 1$

$\Delta\mathcal{L}_{k-1} = -\infty$

for $l = 1, 2, \dots, M_{\text{ls}}$ **do**

$\alpha_l = \epsilon_l / \|\nabla_{\theta}\mathcal{L}(\theta_0)\|_{\mathcal{G}^{-1}(\theta_0)}^2$

$\theta_l = \alpha_l \mathcal{G}^{-1}(\theta_0) \nabla_{\theta}\mathcal{L}(\theta_0)$

$\Delta\mathcal{L}_l = \mathcal{L}_l - \mathcal{L}_0$

if $\Delta\mathcal{L}_l < \Delta\mathcal{L}_{l-1} + \text{tol}_{\Delta\mathcal{L}}$ **then**

return α_{l-1}

end if

$\epsilon_{l+1} = \begin{cases} \eta\epsilon_l & \text{if } \Delta\mathcal{L}_l > \frac{\epsilon_l(1-2\eta)}{2\eta} \\ \frac{\epsilon_l^2}{2(\epsilon_l-\Delta\mathcal{L}_l)} & \text{otherwise} \end{cases}$

end for

E.2 Estimation of the Rényi divergence

In A-POIS, the Rényi divergence needs to be computed between the behavioral $p(\cdot|\theta)$ and target $p(\cdot|\theta')$ distributions on trajectories. This is likely impractical as it requires to integrate over the trajectory space. Moreover, for stochastic environments it cannot be computed unless we know the transition model P . The following result provides an exact, although loose, bound to this quantity in the case of finite-horizon tasks.

Proposition E.1. *Let $p(\cdot|\theta)$ and $p(\cdot|\theta')$ be the behavioral and target trajectory probability density functions. Let $H < \infty$ be the task-horizon. Then, it holds that:*

$$d_{\alpha}(p(\cdot|\theta')\|p(\cdot|\theta)) \leq \left(\sup_{s \in \mathcal{S}} d_{\alpha}(\pi_{\theta'}(\cdot|s)\|\pi_{\theta}(\cdot|s)) \right)^H.$$

Proof. We prove the proposition by induction on the horizon H . We define $d_{\alpha,H}$ as the α -Rényi divergence at horizon H . For $H = 1$ we have:

$$\begin{aligned} d_{\alpha,1}(p(\cdot|\theta')\|p(\cdot|\theta)) &= \int_{\mathcal{S}} D(s_0) \int_{\mathcal{A}} \pi_{\theta}(a_0|s_0) \left(\frac{\pi_{\theta'}(a_0|s_0)}{\pi_{\theta}(a_0|s_0)} \right)^{\alpha} \int_{\mathcal{S}} P(s_1|s_0, a_0) ds_1 da_0 ds_0 \\ &= \int_{\mathcal{S}} D(s_0) \int_{\mathcal{A}} \pi_{\theta}(a_0|s_0) \left(\frac{\pi_{\theta'}(a_0|s_0)}{\pi_{\theta}(a_0|s_0)} \right)^{\alpha} da_0 ds_0 \\ &\leq \int_{\mathcal{S}} D(s_0) ds_0 \sup_{s \in \mathcal{S}} \int_{\mathcal{A}} \pi_{\theta}(a_0|s) \left(\frac{\pi_{\theta'}(a_0|s)}{\pi_{\theta}(a_0|s)} \right)^{\alpha} da_0 \\ &\leq \sup_{s \in \mathcal{S}} d_{\alpha}(\pi_{\theta'}(\cdot|s)\|\pi_{\theta}(\cdot|s)), \end{aligned}$$

where the last but one passage follows from Holder's inequality. Suppose that the proposition holds for any $H' < H$, let us prove the proposition for H .

$$\begin{aligned}
d_{\alpha,H}(p(\cdot|\theta')\|p(\cdot|\theta)) &= \int_{\mathcal{S}} D(s_0) \cdots \int_{\mathcal{A}} \pi_{\theta}(a_{H-2}|s_{H-2}) \left(\frac{\pi_{\theta'}(a_{H-2}|s_{H-2})}{\pi_{\theta}(a_{H-2}|s_{H-2})} \right)^{\alpha} \int_{\mathcal{S}} P(s_{H-1}|s_{H-2}, a_{H-2}) \\
&\quad \times \int_{\mathcal{A}} \pi_{\theta}(a_{H-1}|s_{H-1}) \left(\frac{\pi_{\theta'}(a_{H-1}|s_{H-1})}{\pi_{\theta}(a_{H-1}|s_{H-1})} \right)^{\alpha} \int_{\mathcal{S}} P(s_H|s_{H-1}, a_{H-1}) ds_0 \dots ds_{H-1} \\
&\quad \times da_{H-2} ds_{H-1} da_{H-1} ds_H \\
&= \int_{\mathcal{S}} D(s_0) \cdots \int_{\mathcal{A}} \pi_{\theta}(a_{H-2}|s_{H-2}) \left(\frac{\pi_{\theta'}(a_{H-2}|s_{H-2})}{\pi_{\theta}(a_{H-2}|s_{H-2})} \right)^{\alpha} \int_{\mathcal{S}} P(s_{H-1}|s_{H-2}, a_{H-2}) \\
&\quad \times \int_{\mathcal{A}} \pi_{\theta}(a_{H-1}|s_{H-1}) \left(\frac{\pi_{\theta'}(a_{H-1}|s_{H-1})}{\pi_{\theta}(a_{H-1}|s_{H-1})} \right)^{\alpha} ds_0 \dots ds_{H-1} da_{H-2} ds_{H-1} da_{H-1} \\
&\leq \int_{\mathcal{S}} D(s_0) \cdots \int_{\mathcal{A}} \pi_{\theta}(a_{H-2}|s_{H-2}) \left(\frac{\pi_{\theta'}(a_{H-2}|s_{H-2})}{\pi_{\theta}(a_{H-2}|s_{H-2})} \right)^{\alpha} \int_{\mathcal{S}} P(s_{H-1}|s_{H-2}, a_{H-2}) \\
&\quad \times ds_0 \dots ds_{H-1} da_{H-2} ds_{H-1} \times \sup_{s \in \mathcal{S}} \int_{\mathcal{A}} \pi_{\theta}(a_{H-1}|s) \left(\frac{\pi_{\theta'}(a_{H-1}|s)}{\pi_{\theta}(a_{H-1}|s)} \right)^{\alpha} da_{H-1} \\
&\leq d_{\alpha,H-1}(p(\cdot|\theta')\|p(\cdot|\theta)) \sup_{s \in \mathcal{S}} d_{\alpha}(\pi_{\theta'}(\cdot|s)\|\pi_{\theta}(\cdot|s)) \\
&\leq \left(\sup_{s \in \mathcal{S}} d_{\alpha}(\pi_{\theta'}(\cdot|s)\|\pi_{\theta}(\cdot|s)) \right)^H,
\end{aligned}$$

where we applied Holder's inequality again and the last passage is obtained for the inductive hypothesis. \square

The proposed bound, however, is typically ultraconservative, thus we propose two alternative estimators of the α -Rényi divergence. The first estimator is obtained by simply rephrasing the definition (4) into a sample-based version:

$$\hat{d}_{\alpha}(P\|Q) = \frac{1}{N} \sum_{i=1}^N \left(\frac{p(x_i)}{q(x_i)} \right)^{\alpha} = \frac{1}{N} \sum_{i=1}^N w_{P/Q}^{\alpha}(x_i), \quad (40)$$

where $x_i \sim Q$. This estimator is clearly unbiased and applies to any pair of probability distributions. However, in A-POIS P and Q are distributions over trajectories, their densities are expressed as products, thus the α -Rényi divergence becomes:

$$\begin{aligned}
d_{\alpha}(p(\cdot|\theta')\|p(\cdot|\theta)) &= \int_{\mathcal{T}} p(\cdot|\theta)(\tau) \left(\frac{p(\tau|\theta')}{p(\tau|\theta)} \right)^{\alpha} d\tau = \\
&= \int_{\mathcal{T}} D(s_{\tau,0}) \prod_{t=0}^{H-1} P(s_{\tau,t+1}|s_{\tau,t}, a_{\tau,t}) \prod_{t=0}^{H-1} \pi_{\theta}(a_{\tau,t}|s_{\tau,t}) \left(\frac{\pi_{\theta'}(a_{\tau,t}|s_{\tau,t})}{\pi_{\theta}(a_{\tau,t}|s_{\tau,t})} \right)^{\alpha} d\tau.
\end{aligned}$$

Since both π_{θ} and $\pi_{\theta'}$ are known we are able to compute exactly for each state $d_{\alpha}(\pi_{\theta'}(\cdot|s)\|\pi_{\theta}(\cdot|s))$ with no need to sample the action a . Therefore, we suggest to estimate the Rényi divergence between two trajectory distributions as:

$$\hat{d}_{\alpha}(p(\cdot|\theta')\|p(\cdot|\theta)) = \frac{1}{N} \sum_{i=1}^N \prod_{t=0}^{H-1} d_{\alpha}(\pi_{\theta'}(\cdot|s_{\tau_i,t})\|\pi_{\theta}(\cdot|s_{\tau_i,t})). \quad (41)$$

E.3 Computation of the Fisher Matrix

In A-POIS the Fisher Information Matrix needs to be estimated off-policy from samples. We can use, for this purpose, the IS estimator:

$$\hat{\mathcal{F}}(\theta'/\theta) = \frac{1}{N} \sum_{i=1}^N w_{\theta'/\theta}(\tau_i) \left(\sum_{t=0}^{H-1} \nabla_{\theta'} \log \pi_{\theta'}(a_{\tau_i,t}|s_{\tau_i,t}) \right)^T \left(\sum_{t=0}^{H-1} \nabla_{\theta'} \log \pi_{\theta'}(a_{\tau_i,t}|s_{\tau_i,t}) \right).$$

The SN estimator is obtained by replacing $w_{\theta'/\theta}(\tau_i)$ with $\tilde{w}_{\theta'/\theta}(\tau_i)$. Those estimators become very unreliable when θ' is far from θ , making them difficult to use in practice. On the contrary, in P-POIS

in presence of Gaussian hyperpolicies the FIM can be computed exactly [49]. If the hyperpolicy has diagonal covariance matrix, i.e., $\nu_{\mu, \sigma} = \mathcal{N}(\mu, \text{diag}(\sigma^2))$, the FIM is also diagonal:

$$\mathcal{F}(\mu, \sigma) = \left(\begin{array}{c|c} \text{diag}(1/\sigma^2) & \mathbf{0} \\ \hline \mathbf{0} & 2\mathbf{I} \end{array} \right),$$

where \mathbf{I} is a properly-sized identity matrix.

E.4 Practical surrogate objective functions

In practice, the Rényi divergence term d_2 in the surrogate objective functions presented so far, either exact in P-POIS or approximate in A-POIS, tends to be overly-conservative. To mitigate this problem, by observing that $d_2(P\|Q)/N = 1/\text{ESS}(P\|Q)$ from equation (6) we can replace the whole quantity with an estimator like $\widehat{\text{ESS}}(P\|Q)$, as presented in equation (6). This leads to the following approximated surrogate objective functions:

$$\begin{aligned} \tilde{\mathcal{L}}_{\lambda}^{\text{A-POIS}}(\theta'/\theta) &= \frac{1}{N} \sum_{i=1}^N w_{\theta'/\theta}(\tau_i) R(\tau_i) - \frac{\lambda}{\sqrt{\widehat{\text{ESS}}(p(\cdot|\theta')\|p(\cdot|\theta))}}, \\ \tilde{\mathcal{L}}_{\lambda}^{\text{P-POIS}}(\rho'/\rho) &= \frac{1}{N} \sum_{i=1}^N w_{\rho'/\rho}(\theta_i) R(\tau_i) - \frac{\lambda}{\sqrt{\widehat{\text{ESS}}(\nu_{\rho'}\|\nu_{\rho})}}. \end{aligned}$$

Moreover, in all the experiments, we use the empirical maximum reward in place of the true R_{\max} .

E.5 Practical P-POIS for Deep Neural Policies (N-POIS)

As mentioned in Section 6.2, P-POIS applied to deep neural policies suffers from a curse of dimensionality due to the high number of (scalar) parameters (which are $\sim 10^3$ for the network used in the experiments). The corresponding hyperpolicy is a multi-variate Gaussian (diagonal covariance) with a very high dimensionality. As a result, the Rényi divergence, used as a penalty, is extremely sensitive even to small perturbations, causing an overly-conservative behavior. First, we give up the exact Rényi computation and use the practical surrogate objective function $\tilde{\mathcal{L}}_{\lambda}^{\text{P-POIS}}$ proposed in Appendix E.4. This, however, is not enough. The importance weights, being the products of thousands of probability densities, can easily become zero, preventing any learning. Hence, we decide to group the policy parameters in smaller blocks, and independently learn the corresponding hyperparameters. In general, we can define a family of M orthogonal policy-parameter subspaces $\{\Theta_m \leq \Theta\}_{m=1}^M$, where $V \leq W$ reads “ V is a subspace of W ”. For each Θ_m , we consider a multi-variate diagonal-covariance Gaussian with Θ_m as support, obtaining a corresponding hyperparameter subspace $\mathcal{P}_m \leq \mathcal{P}$. Then, for each \mathcal{P}_m , we compute a separate surrogate objective (where we employ self-normalized importance weights):

$$\tilde{\mathcal{L}}_{\lambda}^{\text{N-POIS}}(\rho'_m/\rho_m) = \frac{1}{N} \sum_{i=1}^N \tilde{w}_{\rho'_m/\rho_m}(\theta_m^i) R(\tau_i) - \frac{\lambda}{\sqrt{\widehat{\text{ESS}}(\nu_{\rho'_m}\|\nu_{\rho_m})}},$$

where $\rho_m, \rho'_m \in \mathcal{P}_m, \theta_m \in \Theta_m$. Each objective is independently optimized via natural gradient ascent, where the step size is found via a line search as usual. It remains to define a meaningful grouping for the policy parameters, i.e., for the weights of the deep neural policy. We choose to group them by network unit, or neuron (counting output units but not input units). More precisely, let denote a network unit as a function:

$$U_i(\mathbf{x}|\theta_m) = g(\mathbf{x}^T \theta_m),$$

where \mathbf{x} is the vector of the inputs to the unit (including a 1 that multiplies the bias parameter) and $g(\cdot)$ is an activation function. To each unit U_m we associate a block Θ_m such that $\theta_m \in \Theta_m$. In more connectivist-friendly terms, we group connections by the neuron they go into. For the network we used in the experiments, this reduces the order of the multivariate Gaussian hyperpolicies from $\sim 10^3$ to $\sim 10^2$. We call this practical variant of our algorithm Neuron-Based POIS (N-POIS). Although some design choices seem rather arbitrary, and independently optimizing hyperparameter blocks clearly neglects some potentially meaningful interactions, the practical results of N-POIS are promising, as reported in Section 6.2. Figure 4 is an ablation study showing the performance of P-POIS variants on Cartpole. Only using both the tricks discussed in this section, we are able to solve the task (this experiment is on 50 iterations only).

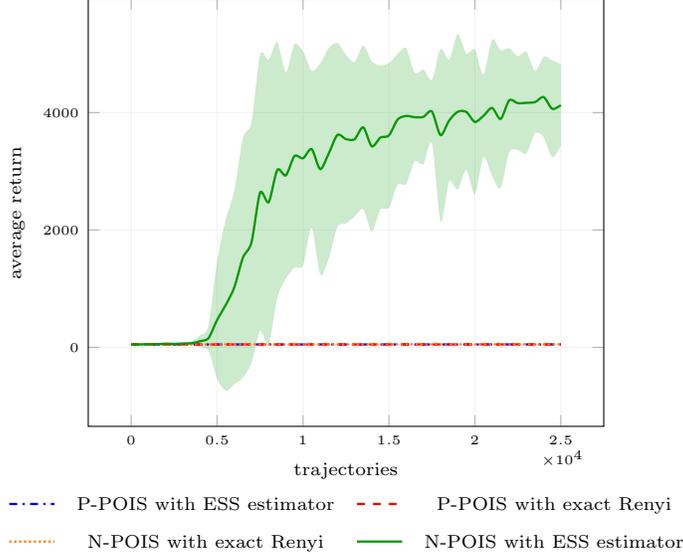


Figure 4: Ablation study for N-POIS (5 runs, 95% c.i.).

F Experiments Details

In this Appendix, we report the hyperparameter values used in the experimental evaluation and some additional plots and experiments. We adopted different criteria to decide the batch size: for linear policies at each iteration 100 episodes are collected regardless of their length, whereas for deep neural policies, in order to be fully comparable with [12], 50000 timesteps are collected at each iteration regardless of the resulting number of episodes (the last episode is cut so that the number of timesteps sums up exactly to 50000). Clearly, this difference is relevant only for episodic tasks.

F.1 Linear policies

In the following we report the hyperparameters shared by all tasks and algorithms for the experiments with linear policies:

- Policy architecture: Normal distribution $\mathcal{N}(u_M(\mathbf{s}), e^{2\Omega})$, where the mean $u_M(\mathbf{s}) = \mathbf{M}\mathbf{s}$ is a linear function in the state variables with no bias, and the variance is state-independent and parametrized as $e^{2\Omega}$, with diagonal Ω .
- Number of runs: 20 (95% c.i.)
- seeds: 10, 109, 904, 160, 570, 662, 963, 100, 746, 236, 247, 689, 153, 947, 307, 42, 950, 315, 545, 178
- Policy initialization: mean parameters sampled from $\mathcal{N}(0, 0.01^2)$, variance initialized to 1
- Task horizon: 500
- Number of iterations: 500
- Maximum number of line search attempts (POIS only): 30
- Maximum number of offline iterations (POIS only): 10
- Episodes per iteration: 100
- Importance weight estimator (POIS only): IS for A-POIS, SN for P-POIS
- Natural gradient (POIS only): No for A-POIS, Yes for P-POIS

Table 3 reports the hyperparameters that have been tuned specifically for each task selecting the best combination based on the runs corresponding to the first 5 seeds.

Table 3: Task-specific hyperparameters for the experiments with linear policy. δ is the significance level for POIS while δ is the step-size for TRPO and PPO. In **bold**, the best hyperparameters found.

Environment	A-POIS (δ)	P-POIS (δ)
Cart-Pole Balancing	0.1, 0.2, 0.3, 0.4 , 0.5	0.1, 0.2, 0.3, 0.4 , 0.5, 0.6, 0.7, 0.8, 0.9 1
Inverted Pendulum	0.8, 0.9 , 0.99, 1	0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8 , 0.9 1
Mountain Car	0.8, 0.9 , 0.99, 1	0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1
Acrobot	0.1, 0.3, 0.5, 0.7 , 0.9	0.1, 0.2 , 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9 1
Double Inverted Pendulum	0.1 , 0.2, 0.3, 0.4, 0.5	0.1 , 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9 1

Environment	TRPO (δ)	PPO (δ)
Cart-Pole Balancing	0.001, 0.01, 0.1 , 1	0.001, 0.01 , 0.1, 1
Inverted Pendulum	0.001, 0.01 , 0.1, 1	0.001, 0.01 , 0.1, 1
Mountain Car	0.001, 0.01 , 0.1, 1	0.001, 0.01, 0.1, 1
Acrobot	0.001, 0.01, 0.1, 1	0.001, 0.01, 0.1, 1
Double Inverted Pendulum	0.001, 0.01, 0.1 , 1	0.001, 0.01, 0.1, 1

F.2 Deep neural policies

In the following we report the hyperparameters shared by all tasks and algorithms for the experiments with deep neural policies:

- Policy architecture: Normal distribution $\mathcal{N}(u_M(\mathbf{s}), e^{2\Omega})$, where the mean $u_M(\mathbf{s})$ is a 3-layers MLP (100, 50, 25) with bias (activation functions: tanh for hidden-layers, linear for output layer), the variance is state-independent and parametrized as $e^{2\Omega}$ with diagonal Ω .
- Number of runs: 5 (95% c.i.)
- seeds: 10, 109, 904, 160, 570
- Policy initialization: uniform Xavier initialization [13], variance initialized to 1
- Task horizon: 500
- Number of iterations: 500
- Maximum number of line search attempts (POIS only): 30
- Maximum number of offline iterations (POIS only): 20
- Timesteps per iteration: 50000
- Importance weight estimator (POIS only): IS for A-POIS, SN for P-POIS
- Natural gradient (POIS only): No for A-POIS, Yes for P-POIS

Table 4 reports the hyperparameters that have been tuned specifically for each task selecting the best combination based on the runs corresponding to the 5 seeds.

Table 4: Task-specific hyperparameters for the experiments with deep neural policies. δ is the significance level for POIS. In **bold**, the best hyperparameters found.

Environment	A-POIS (δ)	P-POIS (δ)
Cart-Pole Balancing	0.9, 0.99 , 0.999	0.4, 0.5, 0.6 , 0.7, 0.8
Mountain Car	0.9, 0.99 , 0.999	0.1, 0.2, 0.3 , 0.4, 0.5, 0.6, 0.7, 0.8
Double Inverted Pendulum	0.9, 0.99 , 0.999	0.4, 0.5, 0.6, 0.7, 0.8
Swimmer	0.9, 0.99 , 0.999	0.4, 0.5, 0.6 , 0.7, 0.8

F.3 Full experimental results

In this section, we report the complete set of results we obtained by testing the two versions of POIS.

In Figure 5 we report additional plots w.r.t. Figure 2 for A-POIS when changing the δ parameter in the Cartpole environment. It is worth noting that the value of δ has also an effect on the speed with which

the variance of the policy approaches zero. Indeed, smaller policy variances induce a larger Rényi divergence and thus with a higher penalization (small δ) reducing the policy variance is discouraged. Moreover, we can see the values of the bound before and after the optimization. Clearly, the higher the value of δ , the higher the value of the bound after the optimization process, as the penalization term is weaker. It is interesting to notice that when $\delta = 1$ the bound after the optimization reaches values that are impossible to reach for any policy and this is a consequence of the high uncertainty in the importance sampling estimator.

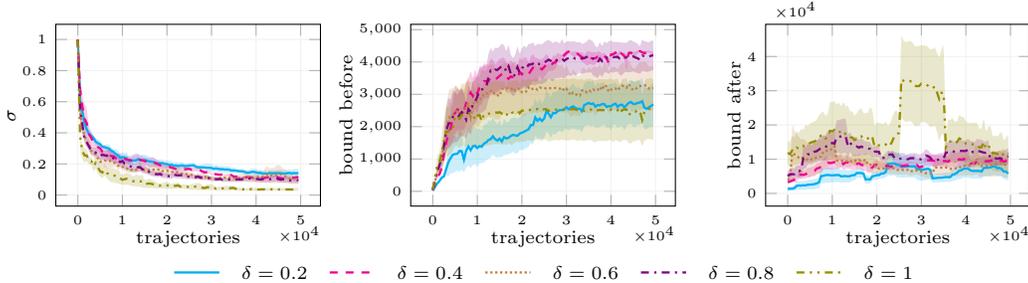


Figure 5: Standard Deviation of the policy (σ), value of the bound before and after the optimization as a function of the number of trajectories for A-POIS in the Cartpole environment for different values of δ (5 runs, 95% c.i.).

We report the comparative table taken from [12] containing all the benchmarked algorithms and the two versions of POIS (Table 5).

Table 5: Cumulative return compared with [12] on deep neural policies (5 runs, 95% c.i.). In **bold**, the performances that are not statistically significantly different from the best algorithm in each task.

Algorithm	Cart-Pole		Double Inverted	
	Balancing	Mountain Car	Pendulum	Swimmer
Random	77.1 \pm 0.0	-415.4 \pm 0.0	149.7 \pm 0.1	-1.7 \pm 0.1
REINFORCE	4693.7 \pm 14.0	-67.1 \pm 1.0	4116.5 \pm 65.2	92.3 \pm 0.1
TNPG	3986.4 \pm 748.9	-66.5 \pm 4.5	4455.4 \pm 37.6	96.0 \pm 0.2
RWR	4861.5 \pm 12.3	-79.4 \pm 1.1	3614.8 \pm 368.1	60.7 \pm 5.5
REPS	565.6 \pm 137.6	-275.6 \pm 166.3	446.7 \pm 114.8	3.8 \pm 3.3
TRPO	4869.8 \pm 37.6	-61.7 \pm 0.9	4412.4 \pm 50.4	96.0 \pm 0.2
DDPG	4634.4 \pm 87.6	-288.4 \pm 170.3	2863.4 \pm 154.0	85.8 \pm 1.8
A-POIS	4842.8 \pm 13.0	-63.7 \pm 0.5	4232.1 \pm 189.5	88.7 \pm 0.55
CEM	4815.4 \pm 4.8	-66.0 \pm 2.4	2566.2 \pm 178.9	68.8 \pm 2.4
CMA-ES	2440.4 \pm 568.3	-85.0 \pm 7.7	1576.1 \pm 51.3	64.9 \pm 1.4
P-POIS	4428.1 \pm 138.6	-78.9 \pm 2.5	3161.4 \pm 959.2	76.8 \pm 1.6

In the following (Figure 6) we show the learning curves of POIS in its two versions for the experiments with deep neural policies.

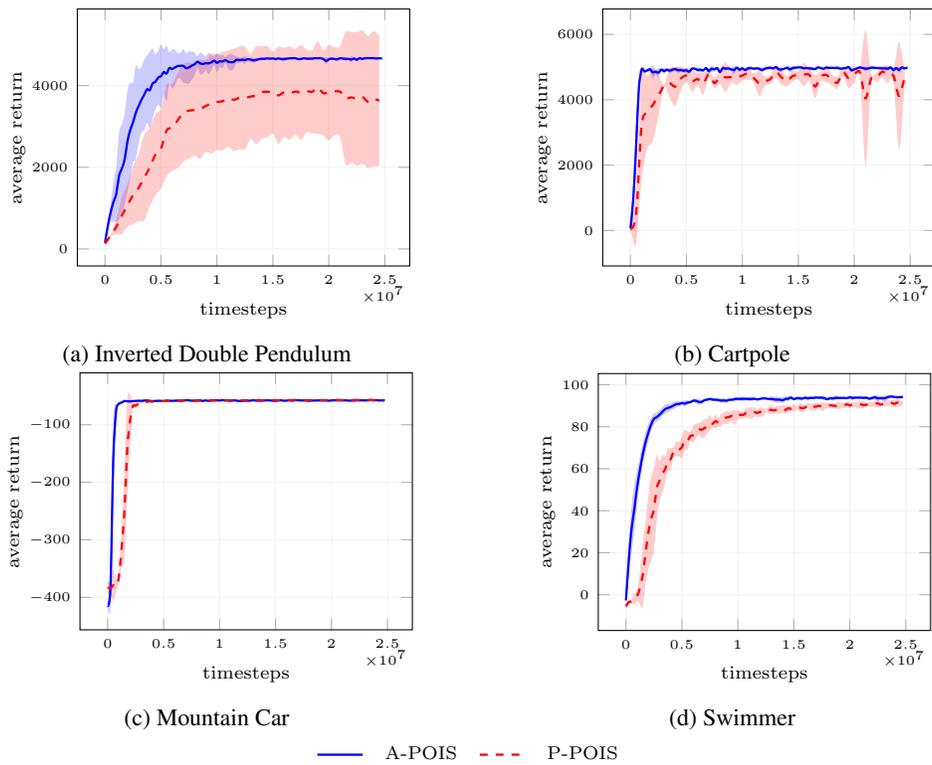


Figure 6: Average return as a function of the number of trajectories for A-POIS, P-POIS with deep neural policies (5 runs, 95% c.i.).