Fingerprint Policy Optimisation
for Robust Reinforcement Learning

Supratik Paul
Department of Computer Science
University of Oxford
supratik.paul@cs.ox.ac.uk

Michael A. Osborne
Department of Engineering Science
University of Oxford
mosb@robots.ox.ac.uk

Shimon Whiteson
Department of Computer Science
University of Oxford
shimon.whiteson@cs.ox.ac.uk

Abstract

Policy gradient methods have been successfully applied to a variety of reinforcement learning tasks. However, while learning in a simulator, these methods do not utilise the opportunity to improve learning by adjusting certain *environment variables*: unobservable state features that are randomly determined by the environment in a physical setting, but that are controllable in a simulator. This can lead to slow learning or convergence to highly suboptimal policies if the environment variable has a large impact on the transition dynamics. In this paper, we present *fingerprint policy optimisation* (FPO) which finds a policy that is optimal in expectation across the distribution of environment variables. The central idea is to use Bayesian optimisation (BO) to actively select the distribution of the environment variable that maximises the improvement generated by each iteration of the policy gradient method. To make this BO practical, we contribute two easy-to-compute low-dimensional *fingerprints* of the current policy. We apply FPO to a number of continuous control tasks of varying difficulty and show that FPO can efficiently learn policies that are robust to significant rare events, which are unlikely to be observable under random sampling but are key to learning good policies.

1   Introduction

Policy gradient methods have demonstrated remarkable success in learning policies for various continuous control tasks Lillicrap et al. (2016); Mordatch et al. (2015); Schulman et al. (2016). However, the expense of running physical trials, coupled with the high sample complexity of these methods, pose significant challenges in directly applying them to a physical setting, e.g., to learn a locomotion policy for a robot. Another problem is evaluating the robustness of a learned policy; it is difficult to ensure that the policy performs as expected, as it is usually infeasible to test it across all possible settings. Fortunately, policies can often be learned and tested in a simulator that exposes key *environment variables* – state features that are unobservable to the agent and randomly determined by the environment in a physical setting, but that are controllable in the simulator.

A naïve application of a policy gradient method updates a policy at each iteration by using a batch of trajectories sampled from the original distribution over environment variables irrespective of the current policy or the training iteration. Thus, it does not explicitly take into account how environment variables affect learning for different policies. Furthermore, this approach is not robust to *significant rare events* (SREs), i.e., it fails any time there are rare events that have significantly different rewards thereby affecting expected performance. Avoiding these problems requires learning *off environment* Frank et al. (2008); Ciosek and Whiteson (2017); Paul et al. (2018): exploiting the ability to adjust environment variables in simulation in order to improve the efficiency and robustness of learning.

For example, consider a quadruped that has to navigate through an environment to reach a goal as quickly as possible. If high velocity policies increase the probability of damage to the quadruped's actuators, then the optimal policy should balance the reward gained by faster locomotion against the cost of potential damage to the robot. The naïve approach is bound to fail in this setting as the low probability of observing a damage, together with the high cost of such an event, yields extremely high variance gradient estimates. However, adjusting environment variables to trigger damage scenarios more often can enable fast learning of robust policies.

In this paper, we propose a new off environment approach called *fingerprint policy optimisation* (FPO) that aims to learn policies that are robust to rare events. At its core, FPO uses a policy gradient method as the policy optimiser. However, unlike the naïve approach, FPO explicitly models the effect of the environment variable on the policy updates, as a function of the policy. Using *Bayesian optimisation* (BO), FPO actively selects the environment distribution at each iteration of the policy gradient method in order to maximise the improvement that one policy gradient update step generates. While this can yield biased gradient estimates, FPO implicitly optimises the bias-variance tradeoff in order to maximise its one-step improvement objective.

A key design challenge in FPO is how to represent the current policy, in cases where the policy is a large neural network with thousands of parameters. To this end, we propose two low-dimensional policy *fingerprints* that act as proxies for the policy. The first approximates the stationary distribution over states induced by the policy, with a size equal to the dimensionality of the state space. The second approximates the policy's marginal distribution over actions, with a size equal to the dimensionality of the action space.

We apply FPO to different continuous control tasks and show that it can outperform existing methods, including those for learning in environments with SREs. Our experiments also show that both fingerprints work equally well in practice, which implies that, for a given problem, the lower dimensional fingerprint can be chosen without sacrificing performance.

2 Problem Setting and Background

A *Markov decision process* (MDP) is a tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}_\theta, r_\theta, s_0, \gamma \rangle$, where $\mathcal{S}$ is the state space, $\mathcal{A}$ the set of actions, $\mathcal{P}$ the transition probabilities, r the reward function, s_0 the probability distribution over the initial state, and $\gamma \in [0, 1)$ the discount factor. We assume that the transition and reward functions depend on some environmental variables θ. At the beginning of each episode, the environment randomly samples θ from some (known) distribution $p(\theta)$. The agent's goal is to learn a policy π mapping states $s \in \mathcal{S}$ to actions $a \in \mathcal{A}$ that maximises the expected return $J(\pi) = \mathbb{E}_\theta[R(\theta, \pi)] = \mathbb{E}_{\theta \sim p(\theta), a_t \sim \pi, s_t, \sim \mathcal{P}}[\sum_t \gamma^t r_\theta(s_t, a_t)]$. With a slight abuse of notation, we use π to denote both the policy and its parameters. We consider environments characterised by significant rare events (SREs), i.e., there exist some low probability values of θ that generate large magnitude returns (positive or negative), yielding a significant impact on $J(\pi)$.

We assume that learning is undertaken in a simulator, or under laboratory conditions where θ can be actively set. This is only mildly restrictive in practice. For example, it is typically possible to expose hidden variables like the coefficient of friction in a simulator, or deliberately limit the power of an actuator to simulate it being damaged.

2.1 Policy Gradient Methods

Starting with some policy π_n at iteration n, gradient based batch policy optimisation methods like REINFORCE Williams (1992), NPG Kakade (2001), and TRPO Schulman et al. (2015) compute an estimate of the gradient $\nabla J(\pi_n)$ by sampling a batch of trajectories from the environment while following policy π_n, and then use this estimate to approximate gradient ascent in J, yielding an updated policy π_{n+1}. REINFORCE uses a fixed learning rate to update the policy; NPG uses the Fisher information matrix to scale the gradient, which makes the updates independent of the policy parameterisation; and TRPO adds a constraint on the KL divergence between consecutive policies.

A major problem for such methods is that the estimate of $\nabla J(\pi)$ can have high variance due to stochasticity in the policy and environment, yielding slow learning Williams (1992); Glynn (1990); Peters and Schaal (2006). In settings with SREs, this problem is compounded by the variance due to θ, which the environment samples for each trajectory in the batch. Furthermore, the SREs may not be

observed during learning since the environment samples $\theta \sim p(\theta)$, which can lead to convergence to a highly suboptimal policy. Since these methods do not explicitly consider the effect of the environment variable on learning, we refer to them as *naïve* approaches.

2.2 Bayesian Optimisation

A *Gaussian process* (GP) Rasmussen and Williams (2005) is a distribution over functions. It is fully specified by its mean $m(\mathbf{x})$ (often assumed to be 0 for convenience) and covariance functions $k(\mathbf{x}, \mathbf{x}')$ which encode any prior belief about the function. A posterior distribution is generated by using observed values to update the belief about the function in a Bayesian way. The squared exponential kernel is a popular choice for the covariance function, and has the form $k(\mathbf{x}, \mathbf{x}') = \sigma_0^2 \exp[-\frac{1}{2}(\mathbf{x} - \mathbf{x}')^T \Sigma^{-1}(\mathbf{x} - \mathbf{x}')]$, where Σ is a diagonal matrix whose diagonal gives lengthscales corresponding to each dimension of $\mathbf{x}$. By conditioning on the observed data, predictions for any new points can be computed analytically as a Gaussian $\mathcal{N}\big(\mu\big(f(\mathbf{x}^*)\big), \sigma^2\big(f(\mathbf{x}^*)\big)\big)$:

$$\mu\big(f(\mathbf{x}^*)\big) = k(\mathbf{x}^*, \mathbf{X})(\mathbf{K} + \sigma_{\text{noise}}^2 \mathbf{I})^{-1} f(\mathbf{X}) \tag{1a}$$

$$\sigma^2\big(f(\mathbf{x}^*)\big) = k(\mathbf{x}^*, \mathbf{x}^*)$$
$$- k(\mathbf{x}^*, \mathbf{X})(\mathbf{K} + \sigma_{\text{noise}}^2 \mathbf{I})^{-1} k(\mathbf{X}, \mathbf{x}^*), \tag{1b}$$

where $\mathbf{X}$ is the vector of observed inputs, $f(\mathbf{X})$ is the corresponding function values, and $\mathbf{K}$ is the covariance matrix with elements $k(\mathbf{x_i}, \mathbf{x_j})$. Probabilistic modelling of the predictions makes GPs well suited for optimising $f(\mathbf{x})$ using BO. Given a set of observations, the next point for evaluation is chosen as the $\mathbf{x}$ that maximises an *acquisition function*, which uses the posterior mean and variance to balance exploitation and exploration. This paper considers two acquisition functions: *upper confidence bound* (UCB) Cox and John (1992, 1997), and *fast information-theoretic Bayesian optimisation* (FITBO) Ru et al. (2017).

Given a dataset $\mathcal{D}_{1:n} = \{(\mathbf{x}_i, f(\mathbf{x}_i))\}_{i=1}^n$, UCB directly incorporates the prediction uncertainty by defining an upper bound: $\alpha_{\text{UCB}}(\mathbf{x} \mid \mathcal{D}_{1:n}) = \mu(f(\mathbf{x}) \mid \mathcal{D}_{1:n}) + \kappa \, \sigma(f(\mathbf{x}) \mid \mathcal{D}_{1:n})$, where κ controls the exploration-exploitation tradeoff. On the other hand, FITBO aims to reduce the uncertainty about the global optimum $f(\mathbf{x}^*)$ by selecting the $\mathbf{x}$ that minimises the entropy of the distribution $p(f(\mathbf{x}^*) \mid \mathbf{x}, \mathcal{D}_{1:n})$. The acquisition function is given by:

$$\alpha_{\text{FITBO}}(\mathbf{x} \mid \mathcal{D}_{1:n}) = H[p(f(\mathbf{x}) \mid \mathbf{x}, \mathcal{D}_{1:n})]$$
$$- \mathbb{E}_{p(f(\mathbf{x}^*)|\mathcal{D}_{1:n})}\big[H[p(f(\mathbf{x}) \mid f(\mathbf{x}^*), \mathbf{x}, \mathcal{D}_{1:n})]\big],$$

which is intractable, but can be computed efficiently following Ru et al. (2017).

BO mainly minimises simple regret: The acquisition function suggests the next point $\mathbf{x}_i$ for evaluation at each timestep, but then the algorithm suggests what it believes to be the optimal point $\hat{\mathbf{x}}_i^*$, and the regret is defined $\sum_{i=1}^N f(\mathbf{x}^*) - f(\hat{\mathbf{x}}_i^*)$. This is different from a bandit setting where the cumulative regret is defined as $\sum_{i=1}^N f(\mathbf{x}^*) - f(\mathbf{x}_i)$. Krause and Ong (2011) showed that the UCB acquisition function is also a viable strategy to minimise cumulative regret in a contextual GP bandit setting, where selection of $\mathbf{x}_i$ conditions on some observed context.

3 Fingerprint Policy Optimisation

To address the challenges posed by environments with SREs, we introduce *fingerprint policy optimisation* (FPO). The main idea is to sample θ from a distribution $q_{\psi_n}(\theta)$ where the parameters ψ_n are conditioned on a fingerprint of the current policy, such that it helps the policy optimisation routine learn a policy that takes into account any SREs. Concretely, FPO executes the following steps at each iteration n. First, it selects ψ_n by approximately solving an optimisation problem defined below. Second, it samples trajectories from the environment using the current policy π_n, where each trajectory uses a values for θ sampled from a distribution parameterised by ψ_n: $\theta \sim q_{\psi_n}(\theta)$. Third, these trajectories are fed to a policy optimisation routine, e.g., a policy gradient algorithm, which uses them to compute an updated policy, $\pi_{n+1} = \text{POLOPT}(\psi, \pi_n)$. Fourth, new, independent trajectories are generated from the environment with π_{n+1} and used to estimate $J(\pi_{n+1})$. Fifth, a new point is added to a dataset $\mathcal{D}_{1:n} = \{\big((\psi_{i-1}, \pi_{i-1}), J(\pi_i)\big)\}_{i=1}^n$, which is input to a GP. The process then repeats, with BO using the GP to select the next ψ.
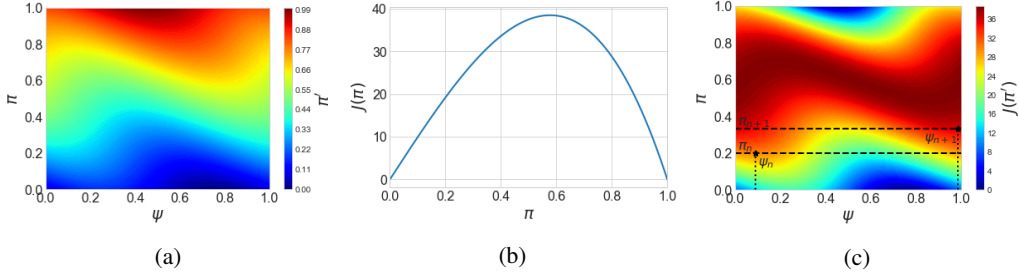
Figure 1: (a) The policy optimisation routine POLOPT takes input (ψ, π) and updates the policy to π'; (b) the return $J(\pi)$ for a policy π; (c) FPO directly models $J(\pi')$ as a function of (ψ, π). At iteration n, π_n is fixed (dashed line) and FPO selects $\psi_n|\pi_n = \text{argmax}_\psi J(\pi_{n+1})$, after which POLOPT$(\psi_n, \pi_n)$ updates the policy to π_{n+1}, and this process repeats.

The key insight behind FPO is that at each iteration n, FPO should select the ψ_n that it expects will maximise the performance of the *next* policy, π_{n+1}:

$$\psi_n|\pi_n = \underset{\psi}{\text{argmax}}\, J(\pi_{n+1})$$
$$= \underset{\psi}{\text{argmax}}\, J\big(\text{POLOPT}(\psi, \pi_n)\big). \tag{2}$$

In other words, FPO chooses the ψ_n that it thinks will help POLOPT maximise the improvement to π_n that can be made in a single policy update. By modelling the relationship between ψ_n, π_n, and $J(\pi_{n+1})$ with a GP, FPO can learn from experience how to select an appropriate ψ for the current π. Modelling $J(\pi_{n+1})$ directly also bypasses the issue of modelling π_{n+1} and its relationship to $J(\pi_{n+1})$, which is infeasible when π_{n+1} is high dimensional. Note that while POLOPT has inputs (ψ, π_n), the optimisation is performed over ψ only, with π_n fixed. Figure 1 illustrates FPO, and is summarised in Algorithm 1. In the remainder of this section, we describe in more detail elements of FPO that are essential to make it work in practice.

3.1 Selecting ψ_n

The optimisation problem in (2) is difficult for two reasons. First, solving it requires calling POLOPT, which is expensive in both computation and samples. Second, the observed $J(\pi)$ can be noisy due to the inherent stochasticity in the policy and the environment.

BO is particularly suited to such settings as it is sample efficient, gradient free, and can work with noisy observations. In this paper, we consider both the UCB and FITBO acquisition functions to select $\psi_n|\pi_n$ in (2), and compare their performance. Formally, we model the returns $J(\pi_n)$ as a GP with two inputs (ψ_{n-1}, π_{n-1}). Given a dataset $\mathcal{D}_{1:n} = \{((\psi_{i-1}, \pi_{i-1}), J(\pi_i))\}_{i=1}^n$, ψ_n is selected by maximising the UCB or FITBO acquisition function:

$$\alpha_{\text{UCB}}(\psi_n|\pi_n) = \mu(J(\pi_{n+1})|\pi_n) + \kappa\sigma(J(\pi_{n+1})|\pi_n) \tag{3}$$
$$\alpha_{\text{FITBO}}(\psi_n|\pi_n) = H[p(J(\pi_{n+1})|\psi, \pi_n)]$$
$$- \mathbb{E}_{p(J(\pi^*_{n+1})|\pi_n)}[H[p(J(\pi_{n+1})|\psi, \pi_n, J(\pi^*_{n+1}))]]. \tag{4}$$

Here we have dropped the $\mathcal{D}_{1:n}$ from all the conditioning sets for ease of notation.

Estimating the gradient estimates using trajectories sampled from the environment with $\theta \sim q_\psi(\theta)$ introduces bias. While most importance sampling (IS) based methods (e.g., Frank et al. (2008) ,Ciosek and Whiteson (2017)) could correct for this bias, FPO does not explicitly do so. Instead FPO lets BO implicitly optimise a bias-variance tradeoff by selecting ψ to maximise the one-step improvement objective.

3.2 Estimating $J(\pi_n)$

Estimating $J(\pi_n)$ accurately in the presence of SREs can be challenging. A Monte Carlo estimate using samples of trajectories from the original environment requires a prohibitively large number of

4

Algorithm 1 Fingerprint Policy Optimisation
---
input Initial policy π_0, original distribution $p(\theta)$, randomly initialised $q_{\psi_0}(\theta)$, policy optimisation
    method POLOPT, number of policy iterations N, dataset $\mathcal{D}_0 = \{\}$
 1: **for** n = 1, 2, 3, . . ., N **do**
 2:    Sample $\theta_{1:k}$ from $q_{\psi_{n-1}}(\theta)$, and with π_{n-1} sample trajectories $\tau_{1:k}$ corresponding to each $\theta_{1:k}$
 3:    Compute $\pi_n = \text{POLOPT}(\tau_{1:k}) = \text{POLOPT}(\psi_{n-1}, \pi_{n-1})$
 4:    Compute $J(\pi_n)$ using numerical quadrature as described in Section 3.2. Use the sampled
       trajectories to compute the policy fingerprint as described in Section 3.3.
 5:    Set $\mathcal{D}_n = \mathcal{D}_{n-1} \cup \{((\psi_{n-1}, \pi_{n-1}), \Delta J(\pi_n))\}$ and update the GP to condition on $\mathcal{D}_n$
 6:    Use either the UCB (3) or FITBO (4) acquisition functions to select ψ_n.
 7: **end for**
---

samples. One alternative would be to apply an IS correction to the trajectories generated from $q_{\psi_n}(\theta)$ for the policy optimisation routine. However, this is not feasible since it would require computing the IS weights $\frac{p(\tau|\theta)}{q_\psi(\tau|\theta)}$, which depend on the unknown transition function. Furthermore, even if the transition function is known, there is no reason why ψ_n should yield a good IS distribution since it is selected with the objective of maximising $J(\pi_{n+1})$.

Instead, FPO applies exhaustive summation for discrete θ and numerical quadrature for continuous θ to estimate $J(\pi_n)$. That is, if the support of θ is discrete, we simply sample a trajectory from each environment defined by θ and estimate $J(\pi_n) = \sum_{l=1}^L p(\theta = \theta_l)R(\theta_l, \pi_n)$. To reduce the variance due to stochasticity in the policy and the environment, we can sample multiple trajectories from each θ_l. Since in practice θ is usually low dimensional, for continuous θ we apply an adaptive Gauss-Kronrod quadrature rule to estimate $J(\pi_n)$. However, we note that due to the limitations of numerical quadrature this approach is unlikely to work for high dimensional θ.

Since for discrete θ we evaluate $J(\pi_n)$ through exhaustive summation, it is natural to consider a variation of the naïve approach, wherein $\nabla J(\pi_n)$ is also evaluated in the same manner during training, i.e. $\nabla J(\pi_n) = \sum_{l=1}^L p(\theta = \theta_l)\nabla J(\theta_l, \pi_n)$. We call this the 'Enum' baseline since the gradient is estimated by enumerating over all possible values of θ. Our experiments in Section 5 show that this baseline is unable to match the performance of FPO.

3.3 Policy Fingerprints

True global optimisation is limited by the curse of dimensionality to low-dimensional inputs, and BO has had only rare successes in problems with more than twenty dimensions Wang et al. (2013). In FPO, many of the inputs to the GP are policy parameters: in practice, the policy is a neural network that may have thousands of parameters, far too many for a GP. Thus, we need to develop a policy *fingerprint*, i.e., a representation that is low dimensional enough to be treated as an input to the GP but expressive enough to distinguish the policy from others.

Foerster et al. (2017) showed that a surprisingly simple fingerprint, consisting only of the training iteration, suffices to stabilise multi-agent Q-learning. Using the training iteration alone as the fingerprint proves insufficient for FPO, as the GP fails to model the response surface and treats all observed $J(\pi)$ as noise. However, the principle still applies: a simplistic fingerprint that discards much information about the policy can still provide sufficient for decision making, in this case to select ψ_n.

In this spirit, we propose two fingerprints. The first, the *state fingerprint*, augments the training iteration with an estimate of the stationary state distribution induced by the policy. In particular, we fit an anisotropic Gaussian to the set of states visited in the trajectories sampled while estimating $J(\pi)$ (see Section 3.2). The size of this fingerprint grows linearly with the dimensionality of the state space, instead of the number of parameters in the policy.

In many settings, the state space is high dimensional, but the action space is low dimensional. Therefore, our second fingerprint, the *action fingerprint*, is a Gaussian approximation of the marginal distribution over actions induced by the policy: $a(\pi) = \int \pi(a|s)s(\pi)ds$ (here $s(\pi)$ is the stationary state distribution induced by π), sampled from trajectories as with the state fingerprint.

Of course, neither the stationary state distribution nor the marginal action distribution are likely to be Gaussian and could in fact be multimodal. Furthermore, the state distribution is estimated from samples used to estimate $J(\pi)$, and not from $p(\theta)$. However, as our results show, these representations are nonetheless effective, as they do not need to accurately describe each policy, but instead just serve as low dimensional fingerprints on which FPO conditions.

3.4 Covariance Function

Our choice of policy fingerprints means that one of the inputs to the GP is a probability distribution. Thus for our GP prior we use a product of three covariance functions, $k_1(\psi, \psi') \, k_2(n, n') \, k_3\big(\text{FGP}(\pi), \text{FGP}(\pi')\big)$, where each of k_1, k_2, and k_3 is a squared exponential covariance function and $\text{FGP}(\pi)$ is the state or action fingerprint of π. Similar to Malkomes et al. (2016), we use the Hellinger distance to replace the Euclidean in k_3: this covariance remains positive-semi-definite as the Hellinger is effectively a modified Euclidean.

4 Related Work

Various methods have been proposed for learning in the presence of SREs. These are usually off environment and either based on learning a good IS distribution from which to sample the environment variable Frank et al. (2008); Ciosek and Whiteson (2017), or Bayesian active selection of the environment variable during learning Paul et al. (2018).

Frank et al. (2008) propose a temporal difference based method that uses IS for efficiently evaluating policies whose expected value may be substantially affected by rare events. However, their method assumes prior knowledge of the SREs, such that they can directly alter the probability of such events during policy evaluation. By contrast, FPO does not require any such prior knowledge about SREs, or the environment variable settings that might trigger them. It only assumes that the original distribution of the environment variable is known, and that the environment variable is controllable during learning.

OFFER Ciosek and Whiteson (2017) is a policy gradient method based algorithm that uses observed trials to gradually changes the IS distribution over the environment variable. Like FPO, it makes no prior assumptions about SREs. However, at each iteration it updates the environment distribution with the objective of minimising the variance of the gradient estimate, which may not lead to the distribution that optimises the learning of the policy. A further disadvantage of the method is that it requires a full transition model of the environment to compute the IS weights. It can also lead to unstable IS estimates if the environment variable affects any transitions besides the initial state.

ALOQ Paul et al. (2018) is a Bayesian optimisation and quadrature based method that models the return as a GP with the policy parameters and environment variable as inputs. At each iteration it actively selects the policy and then the environment variable in an alternating fashion and as such, performs the policy search in the parameter space. Being a BO based method, it does not make any assumption of the Markov property of the environment and is highly sample efficient. However, it can only be applied to settings with low dimensional policies. Furthermore, its computational cost scales cubically with the number of iterations, and is thus limited to settings where a good policy can be found within relatively few iterations. In contrast to this, FPO uses a policy gradient method to perform policy optimisation while the BO component generates trajectories, which when used by the policy optimiser are expected to lead to a larger improvement in the policy.

In the wider BO literature, Williams et al. (2000) suggested a method for settings where the objective is to optimise expensive integrands. However, their method does not specifically consider the impact of SREs and, as shown by Paul et al. (2018), are unsuitable for such settings. Toscano-Palmerin and Frazier (2018) suggest BQO, another BO based method for expensive integrands. Their method also does not explicitly consider SREs. Finally, both these methods suffer from all the disadvantages of BO based methods mentioned earlier.

Rajeswaran et al. (2017) propose a different, off-environment approach for cases where the simulator settings can have a significant impact on policy performance. Their algorithm, EPOpt(ϵ), seeks to learn robust policies by maximising the ϵ-percentile conditional value at risk (CVaR) of the policy. First, it randomly samples a set of simulator settings; then trajectories are sampled for each of these settings. A policy optimisation routine (e.g., TRPO Schulman et al. (2015)) is then used to update

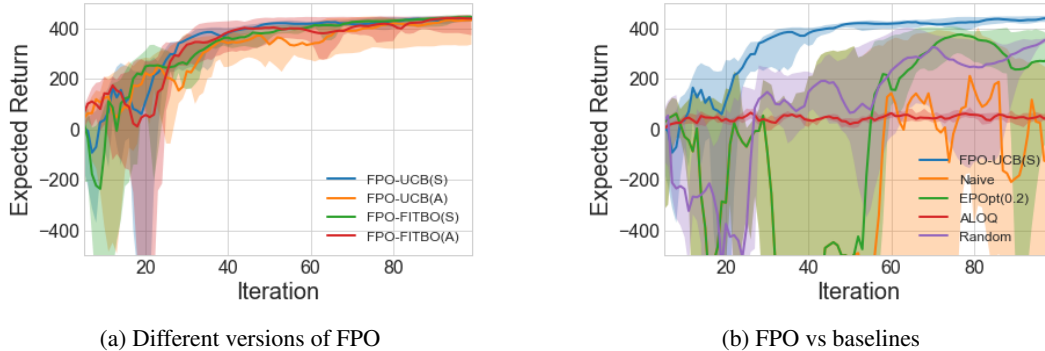(a) Different versions of FPO   (b) FPO vs baselines

Figure 2: Performance of FPO on the Cliff Walker problem. The expected return shown is smoothed by averaging across the last 5 iterations.

the policy based on only those trajectories with returns lower than the ϵ percentile in the batch. A fundamental difference to FPO is that it finds a risk-averse solution based on CVaR, while FPO finds a risk neutral policy. Also, while FPO actively changes the distribution for sampling the environment variable at each iteration, EPOpt samples them from the original distribution, and is thus unlikely to be suitable for settings with SREs, since it will not generate them often enough to to learn an appropriate response. Finally, EPOpt discards all sampled trajectories in the batch with returns greater than ϵ percentile for use by the policy optimisation routine, making it highly sample inefficient, especially for low values of ϵ.

RARL proposed by Pinto et al. (2017) also seeks to learn robust policies by training in a simulator where an adversary applies destabilising forces with both the agent and the adversary being trained simultaneously. RARL requires significant prior knowledge in setting up the adversary to ensure that it strikes a balance between making the environment so difficult that the agent in unable to learn, and making it so easy that the policy learnt by the agent isn't robust enough. Also, like EPOpt it does not consider any settings with SREs.

By learning an optimal distribution for the environment variable conditioned on the policy fingerprint, FPO also has some parallels with meta-learning. Methods like MAML Finn et al. (2017), and Reptile Nichol et al. (2018) seek to find a good policy representation that can be adapted quickly to a specified task. Andrychowicz et al. (2016); Chen et al. (2017) seek to optimise neural networks by learning an automatic update rule based on transferring knowledge from similar optimisation problems. To maximise the performance of a neural network across a set of discrete tasks, Graves et al. (2017) propose a method for automatically selecting a curriculum during learning. Their method treats the problem as a multi-armed bandit and use the Exp3 algorithm Auer et al. (2002) to find the optimal curriculum. Unlike these methods, which seek to quickly adapt to a new task after training on some related task, FPO seeks to maximise the expected return across a variety of tasks.

5 Experiments

To evaluate the empirical performance of FPO, we start by applying it to a simple problem: a modified version of the cliff walker task Sutton and Barto (1998), with one dimensional state and action spaces. We then move on to simulated robotics problems based on the MuJoCo simulator Brockman et al. (2016) with much higher dimensionalities. These were modified to include SREs. We aim to answer two questions: (1) How do the different versions of FPO (UCB vs. FITBO acquisition functions, state (S) vs. action (A) fingerprints) compare with each other? (2) How does FPO compare to existing methods (Naïve, Enum, OFFER, EPOpt, ALOQ), and ablated versions of FPO. We repeat all our experiments across 10 random starts, and present the median (solid line) and quartiles (shaded region) in the plots. We use TRPO as the policy optimisation method, and neural net policies of (5,5) hidden units for the Cliff Walker experiment, and (100,100) hidden units with relu activations for the Half Cheetah and Ant experiments. We used rllab Duan et al. (2016) as the codebase for TRPO, and all other hyperparameters were kept to their default values.

(a) Half Cheetah task

(b) Different versions of FPO

(c) FPO vs baselines
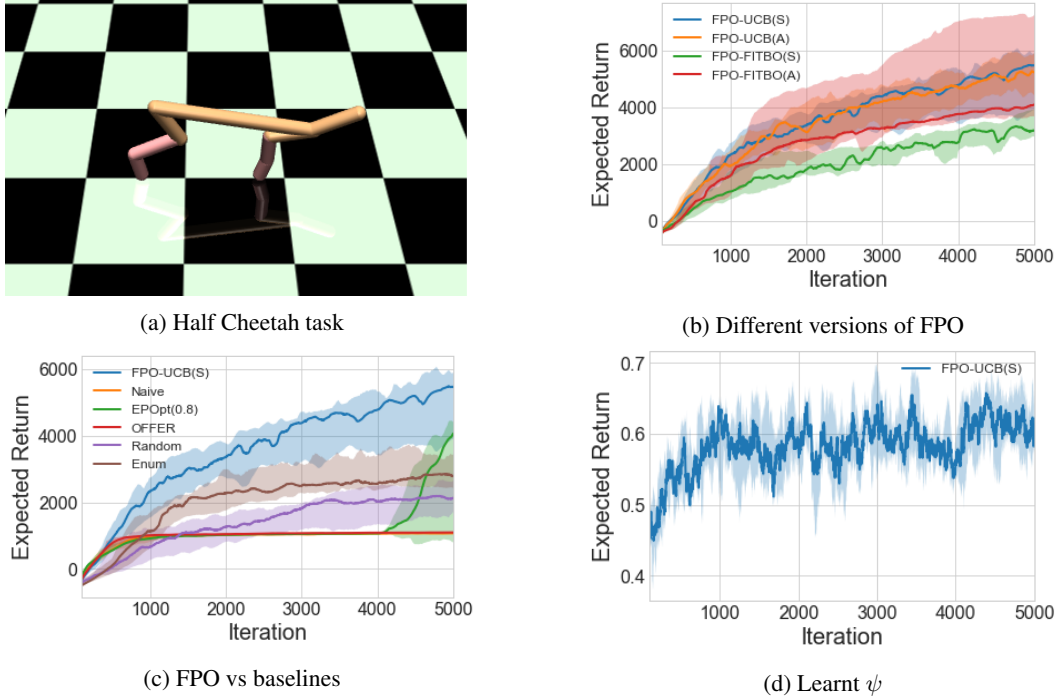
(d) Learnt ψ

Figure 3: The half cheetah task and comparison of performance of FPO. The expected return shown is smoothed by averaging across the last 100 iterations. We also show the ψ (in this case probability of the velocity target being 4) selected by FPO across iterations.

Due to the disadvantages of ALOQ mentioned in Section 4, we were able to apply it only on the cliff walker problem. The policy dimensionality and the total number of iterations for the simulated robotic tasks were far too high. Note also that, while we compare FPO to EPOpt, these methods optimise for different objectives.

5.1 Cliff Walker

We start with a toy problem: a modified version of the cliff walker problem where instead of a gridworld we consider an agent moving in a continuous state space $\mathbb{R}$; the agent starts randomly near the state 0, and at each timestep can take an action $a \in \mathbb{R}$. The environment then transitions the agent to a new location $s' = s + 0.025\mathrm{sign}(a) + 0.005\epsilon$, where ϵ is standard Gaussian noise. The location of the cliff is given by $1 + \theta$, where θ follows a Beta distribution. If the agent's current state is lower than the cliff location, it gets a reward equal to its state; otherwise it falls off the cliff and gets a reward of -5000, terminating the episode. Thus the objective is to learn to walk as close to the cliff edge as possible without falling over.

We can see from Figure 2a that all versions of FPO do equally well on the task. However, in Figure 2b we see that not only does FPO-UCB(S) learn a policy with a higher expected return, it also has a much lower variance than all other baselines, except for ALOQ. This is not surprising since, as discussed in Section 2.1, the gradient estimates without active selection of ψ are likely to have high variance due to the presence of SREs. For EPOpt we set $\epsilon = 0.2$ and performed rejection sampling after 50 iterations. The poor performance of ALOQ is expected since even in this simple problem, the policy dimensionality is 47, which is still quite high for a BO based method. We could not run OFFER since an analytical solution of $\nabla_\psi q_\psi(\theta)$ does not exist.

5.2 Half Cheetah

Next we consider simulated robotic locomotion tasks using the Mujoco simulator. In the half cheetah task shown in Figure 3a, the objective is to learn a locomotion policy for the planar robot with two legs. We modified the original problem such that in 98% of the cases the objective of the agent is to

(a) Ant task

(b) Different versions of FPO

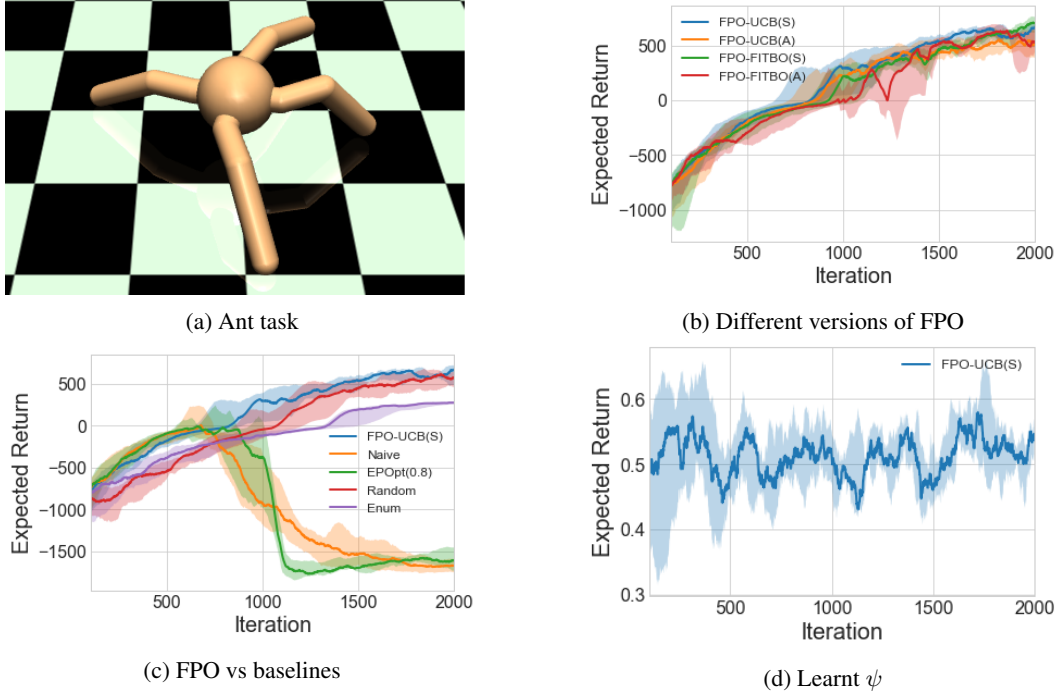
(c) FPO vs baselines

(d) Learnt ψ

Figure 4: The ant task and comparison of performance of FPO. The expected return shown is smoothed by averaging across the last 100 iterations. We also show the ψ (in this case probability of damage to the ant if the velocity is greater than 2) selected by FPO across iterations.

achieve a target velocity of 2, with the rewards decreasing linearly for being away from the target. In the remaining 2%, the target velocity is set to 4, with a large bonus reward, which acts as an SRE.

Figure 3b shows that FPO with the UCB acquisition function outperforms the FITBO acquisition function. We suspect that this is because FITBO tends to over-explore. Also, as mentioned in Section 2.2, it was developed with the aim of minimising simple regret and it is not known how efficient it is at minimising cumulative regret. Like in the cliff walker experiment, both the action and state fingerprints perform equally well with the UCB acquisition.

Figure 3c shows that the Naïve method and OFFER converge to a locally optimal policy that completely ignores the SRE, while the random selection of ψ does slightly better. The Enum baseline performs even better than the random baseline, but is still far worse than FPO. This shows that the key to learning a good policy is the active selection of ψ_n, which also includes the implicit bias-variance tradeoff as performed by FPO. We set $\epsilon = 0.8$ for EPOpt, but in this case we use the trajectories with returns exceeding the threshold for the policy optimisation since the SRE has a large positive return. Although its performance increases after iteration 4,000, it is extremely sample inefficient, requiring about five times the samples of FPO. We did not run ALOQ as it is entirely infeasible given the policy dimensionality is more than 20,000.

Finally, in 3d we present the schedule of ψ as selected by FPO-UCB(S) across different iterations. We see that FPO manages to learn to vary ψ, starting with 0.5 initially, to hovering around 0.6 once the cheetah is able to consistently reach velocities greater than 2.

5.3 Ant

The ant environment shown in Figure 4a is a much more difficult problem than the half cheetah since the agent now moves in 3D, and the state space has 111 dimensions compared to 17 for half cheetah. The larger state space also makes learning difficult due to the higher variance in the gradient estimates. We modified the original problem such that velocities greater than 2 carries a 5% chance of damage to the ant. On incurring damage, which we treat as the SRE, the agent receives a large negative reward, and the episode terminates.

9

In Figure 4b we compare the performance of the UCB versions of FPO. We see that there is no significant difference between the performance of the state or action fingerprint, or between the UCB and FITBO acquisition functions. Thus we see that the lower dimensional fingerprint (in this case the action fingerprint) can be chosen without any significant drop in performance.

In Figure 4c we see that for the Naïve method and EPOpt the performances drops significantly after about 750 iterations. This is because the policies being learned by these methods lead to velocities beyond 2, which after factoring in the effect of the SRE, leads to much lower expected returns. Since the SREs are not seen often enough, these methods do not learn that higher velocities actually lead to lower expected returns. The Enum baseline once again performs better than the naïve approach, but is still unable to match the performance of FPO. The random baseline tends to perform better, and eventually matches the performance of FPO. We could not run OFFER in this setting as computing the IS weights in this setting would require knowledge of the transition model.

In Figure 4d we see that the optimal schedule for ψ as learnt by FPO-UCB(S) tends to hover around 0.5. This goes toward explaining relatively good performance of the random baseline: since the baseline samples $\psi \sim U(0, 1)$ at each iteration, the expected value of ψ is 0.5, and thus we can expect it to find a good policy eventually. Of course the active selection of ψ still matters, and this is borne out by FPO outperforming it initially.

6   Conclusion

In this paper we presented FPO, a method based on the insight that active selection of the environment variable during learning can lead to policies that take into account the effect of SREs. We introduced novel state and action fingerprints that can be used by BO with a one-step improvement objective to make FPO scalable to high dimensional tasks irrespective of the policy dimensionality. We applied FPO to a number of continuous control tasks of varying difficulty and showed that FPO can efficiently learn policies that are robust to significant rare events, which are unlikely to be observable under random sampling but are key to learning good policies. In the future we would like to develop fingerprints for discrete state and action spaces, and explore using a multi-step improvement objective for the BO component.

Acknowledgements

References

Andrychowicz, M., Denil, M., Gómez, S., Hoffman, M. W., Pfau, D., Schaul, T., and de Freitas, N. (2016). Learning to learn by gradient descent by gradient descent. In *Neural Information Processing Systems (NIPS)*.

Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. (2002). The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77.

Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. (2016). Openai gym.

Chen, Y., Hoffman, M. W., Colmenarejo, S. G., Denil, M., Lillicrap, T. P., Botvinick, M., and de Freitas, N. (2017). Learning to learn without gradient descent by gradient descent. In *International Conference on Machine Learning (ICML)*.

Ciosek, K. and Whiteson, S. (2017). Offer: Off-environment reinforcement learning. In *AAAI Conference on Artificial Intelligence*.

Cox, D. D. and John, S. (1992). A statistical method for global optimization. In *IEEE International Conference on Systems, Man and Cybernetics*.

Cox, D. D. and John, S. (1997). SDO: A statistical method for global optimization. In *in Multidisciplinary Design Optimization: State-of-the-Art*, pages 315–329.

Duan, Y., Chen, X., Houthooft, R., Schulman, J., and Abbeel, P. (2016). Benchmarking deep reinforcement learning for continuous control. In *International Conference on Machine Learning (ICML)*.

Finn, C., Abbeel, P., and Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning (ICML)*.

Foerster, J., Nardelli, N., Farquhar, G., Torr, P., Kohli, P., and Whiteson, S. (2017). Stabilising experience replay for deep multi-agent reinforcement learning. In *International Conference on Machine Learning (ICML)*.

Frank, J., Mannor, S., and Precup, D. (2008). Reinforcement learning in the presence of rare events. In *International Conference on Machine Learning (ICML)*.

Glynn, P. W. (1990). Likelihood ratio gradient estimation for stochastic systems. In *Communications of the ACM*.

Graves, A., Bellemare, M. G., Menick, J., Munos, R., and Kavukcuoglu, K. (2017). Automated curriculum learning for neural networks. In *International Conference on Machine Learning (ICML)*.

Kakade, S. (2001). A natural policy gradient. In *Neural Information Processing Systems (NIPS)*.

Krause, A. and Ong, C. S. (2011). Contextual gaussian process bandit optimization. In *Neural Information Processing Systems (NIPS)*.

Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. (2016). Continuous control with deep reinforcement learning. In *International Conference on Learning Representations (ICLR)*.

Malkomes, G., Schaff, C., and Garnett, R. (2016). Bayesian optimization for automated model selection. In *Neural Information Processing Systems (NIPS)*.

Mordatch, I., Lowrey, K., Andrew, G., Popovic, Z., and Todorov, E. V. (2015). Interactive control of diverse complex characters with neural networks. In *Neural Information Processing Systems (NIPS)*.

Nichol, A., Achiam, J., and Schulman, J. (2018). On first-order meta-learning algorithms. *CoRR*, abs/1803.02999.

Paul, S., Chatzilygeroudis, K., Ciosek, K., Mouret, J.-B., Osborne, M., and Whiteson, S. (2018). Alternating optimisation and quadrature for robust control. In *AAAI Conference on Artificial Intelligence*.

Peters, J. and Schaal, S. (2006). Policy gradient methods for robotics. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*.

Pinto, L., Davidson, J., Sukthankar, R., and Gupta, A. (2017). Robust adversarial reinforcement learning. In *International Conference on Machine Learning (ICML)*.

Rajeswaran, A., Ghotra, S., Levine, S., and Ravindran, B. (2017). EPOpt: Learning robust neural network policies using model ensembles. *International Conference on Learning Representations (ICLR)*.

Rasmussen, C. E. and Williams, C. K. I. (2005). *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press.

Ru, B., McLeod, M., Granziol, D., and Osborne, M. A. (2017). Fast Information-theoretic Bayesian Optimisation. In *International Conference on Machine Learning (ICML)*.

Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. (2015). Trust region policy optimization. In *International Conference on Machine Learning (ICML)*.

Schulman, J., Moritz, P., Levine, S., Jordan, M., and Abbeel, P. (2016). High-dimensional continuous control using generalized advantage estimation. In *Proceedings of the International Conference on Learning Representations (ICLR)*.

Sutton, R. S. and Barto, A. G. (1998). *Reinforcement Learning : An Introduction*. MIT Press.

Toscano-Palmerin, S. and Frazier, P. I. (2018). Bayesian Optimization with Expensive Integrands. *ArXiv e-prints*.

Wang, Z., Zoghi, M., Hutter, F., Matheson, D., and De Freitas, N. (2013). Bayesian Optimization in High Dimensions via Random Embeddings. In *IJCAI*, pages 1778–1784.

Williams, B. J., Santner, T. J., and Notz, W. I. (2000). Sequential design of computer experiments to minimize integrated response functions. *Statistica Sinica*.

Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*.