

# Ensemble Policy Optimization for Learning Robust Policies

Aravind Rajeswaran<sup>1</sup>, Sarvjeet Ghotra<sup>2</sup>, Balaraman Ravindran<sup>3</sup>, Sergey Levine<sup>4</sup>

<sup>1</sup> UW Seattle, <sup>2</sup> NITK Surathkal, <sup>3</sup> IIT Madras, <sup>4</sup> UC Berkeley

aravraj@cs.washington.edu, sarvjeet.13it236@nitk.edu.in, ravi@cse.iitm.ac.in, svlevine@eecs.berkeley.edu

**Abstract**—Sample complexity and safety are major challenges when learning policies with reinforcement learning for real-world tasks, especially when the policies are represented using rich function approximators like deep neural networks. Model-based methods where the real-world target domain is approximated using a simulated source domain provide an avenue to tackle the above challenges by augmenting real data with simulated data. However, discrepancies between the simulated source domain and the target domain pose a challenge for simulated training. We introduce the EPOpt algorithm, which uses an ensemble of simulated source domains and a form of adversarial training to learn policies that are robust and generalize to a broad range of possible target domains, including unmodeled effects. Further, the probability distribution over source domains in the ensemble can be adapted using data from target domain and approximate Bayesian methods, to progressively make it a better approximation. Thus, learning on a model ensemble, along with source domain adaptation, provides the benefit of both robustness and learning/adaptation.

## I. INTRODUCTION

Reinforcement learning, both with shallow function approximators, as well as multi-layer architectures (deep RL), has recently demonstrated remarkable success in a wide range of tasks like games [19, 30, 15], simulated control problems [16, 21, 26], and graphics [23]. However, high sample complexity is a major barrier for directly applying model-free deep RL methods for physical control tasks. Model-free algorithms like Q-learning, actor-critic, and policy gradients are known to suffer from long learning times [12], which is compounded when used in conjunction with expressive function approximators like deep neural networks (DNNs). The challenge of gathering samples from the real world is further exacerbated by issues of safety for the agent and environment, since sampling with partially learned policies could be unstable [9]. Thus, model-free deep RL methods often require a prohibitively large numbers of potentially dangerous samples for physical control tasks.

Model-based methods, where the real-world target domain is approximated with a simulated source domain, provide an avenue to tackle the above challenges by learning policies using simulated data. The principal challenge with simulated training is the systematic discrepancy between source and target domains, and therefore, methods that compensate for systematic discrepancies (modeling errors) are needed to transfer results from simulations to real world using RL. We show that the impact of such discrepancies can be mitigated through two key ideas: (1) training on an ensemble of models in an adversarial fashion to learn policies that are robust to parametric model errors, as well as to unmodeled effects; and (2) adaptation of

the source domain ensemble using data from the target domain to progressively make it a better approximation. This can be viewed either as an instance of model-based Bayesian RL [10]; or as transfer learning from a collection of simulated source domains to a real-world target domain [33]. While a number of model-free RL algorithms have been proposed (see, e.g., [6] for a survey), their high sample complexity demands use of a simulator, effectively making them model-based. We show in our experiments that such methods learn policies which are highly optimized for the specific models used in the simulator, but are brittle under model mismatch. This is not surprising, since deep networks are remarkably proficient at exploiting any systematic regularities in a simulator. Recent work have also demonstrated similar brittleness to perturbations in the state distribution, induced for example through application of a short lived external force [26]. Addressing robustness of DNN-policies is particularly important to transfer their success from simulated tasks to physical systems.

In this paper, we propose the Ensemble Policy Optimization (EPOpt- $\epsilon$ ) algorithm for finding policies that are robust to model mismatch. In line with model-based Bayesian RL, we learn a policy for the target domain by alternating between two phases: (i) given a source (model) distribution (i.e. ensemble of models), find a robust policy that is competent for the whole distribution; (ii) gather data from the target domain using said robust policy, and adapt the source distribution. EPOpt uses an ensemble of models sampled from the source distribution, and a form of adversarial training to learn robust policies that generalize to a broad range of models. By robust, we mean insensitivity to parametric model errors and broadly competent performance for *direct-transfer* (also referred to as *jumpstart* like in [33]). Direct-transfer performance refers to the average initial performance (return) in the target domain, without any direct training on the target domain. By adversarial training, we mean that model instances on which the policy performs poorly in the source distribution are sampled more often in order to encourage learning of policies that perform well for a wide range of model instances. This is in contrast to methods which learn highly optimized policies for specific model instances, but brittle under model perturbations. In our experiments, we did not observe significant loss in performance by requiring the policy to work on multiple models (for example, through adopting a more conservative strategy). Further, we show that policies learned using EPOpt are robust even to effects not modeled in the source domain. Such unmodeled effects are a major issue when transferring from simulation to the real world. For the model adaptation step (ii), we present a simple method

using approximate Bayesian updates, which progressively makes the source distribution a better approximation of the target domain. As a proof of concept, we demonstrate the proposed methods on the hopper benchmark in MuJoCo. Our experimental results suggest that adversarial training on model ensembles produces robust policies which generalize better than policies trained on a single, maximum-likelihood model (of source distribution) alone.

## II. PROBLEM FORMULATION

We consider parametrized Markov Decision Processes (MDPs), which are tuples of the form:  $\mathcal{M}(p) \equiv \langle \mathcal{S}, \mathcal{A}, \mathcal{T}_p, \mathcal{R}_p, \gamma, S_{0,p} \rangle$  where  $\mathcal{S}$ ,  $\mathcal{A}$  are (continuous) states and actions respectively;  $\mathcal{T}_p$ ,  $\mathcal{R}_p$ , and  $S_{0,p}$  are the state transition, reward function, and initial state distribution respectively, all parametrized by  $p$ ; and  $\gamma$  is the discount factor. Thus, we consider a set of MDPs with the same state and action spaces. Each MDP in this set could potentially have different transition functions, rewards, and initial state distributions. We use transition functions of the form  $S_{t+1} \equiv \mathcal{T}_p(s_t, a_t)$  where  $\mathcal{T}_p$  is a random process and  $S_{t+1}$  is a random variable.

We distinguish between source and target MDPs using  $\mathcal{M}$  and  $\mathcal{W}$  respectively, and also refer to them as source and target *domain*, as is common in the transfer learning set-up. Our objective is to learn the optimal policy for  $\mathcal{W}$ ; and to do so, we have access to  $\mathcal{M}(p)$ . We assume that we have a distribution ( $\mathcal{D}$ ) over the source domains (MDPs) generated by a distribution over the parameters  $P \equiv \mathcal{P}(p)$  that capture our subjective belief about the parameters of  $\mathcal{W}$ . Let  $\mathcal{P}$  be parametrized by  $\psi$  (e.g. mean, standard deviation). For example,  $\mathcal{M}$  could be a hopping task with reward proportional to hopping velocity and falling down corresponds to a terminal state. For this task,  $p$  could correspond to parameters like torso mass, ground friction, and damping in joints, all of which affect the dynamics. Ideally, we would like the target domain to be in the model class, i.e.  $\{\exists p \mid \mathcal{M}(p) = \mathcal{W}\}$ . However, in practice, there are likely to be unmodeled effects, and we analyze this setting in our experiments. We wish to learn a policy  $\pi_\theta^*(s)$  that performs well for all  $\mathcal{M} \sim \mathcal{D}$ . Note that this robust policy does not have an explicit dependence on  $p$ , and we require it to perform well without knowledge of  $p$ .

## III. LEARNING PROTOCOL AND EPOPT ALGORITHM

We follow the round-based learning protocol of Bayesian model-based RL. We use the term *rounds* when interacting with the target domain, and *episode* when performing rollouts with the simulator. In each round, we interact with the target domain after computing the robust policy on the current (i.e. posterior) simulated source distribution. Following this, we update the source distribution using data from the target domain collected by executing the robust policy. Thus, in round  $i$ , we update two sets of parameters:  $\theta_i$ , the parameters of the robust policy (neural network); and  $\psi_i$ , the parameters of the source distribution. The two key steps in this procedure are finding a robust policy given a source distribution; and updating the

source distribution using data from the target domain. In this section, we present our approach for both of these steps.

### A. Robust policy search

We introduce the EPOpt algorithm for finding a robust policy using the source distribution. EPOpt is a policy gradient based meta-algorithm which uses batch policy optimization methods as a subroutine. Batch policy optimization algorithms [38, 11, 29] collect a batch of trajectories by rolling out the current policy, and use the trajectories to make a policy update. The basic structure of EPOpt is to sample a collection of models from the source distribution, sample trajectories from each of these models, and make a gradient update based on a subset of sampled trajectories. We first define evaluation metrics for the parametrized policy,  $\pi_\theta$ :

$$\begin{aligned} \eta_{\mathcal{M}}(\theta, p) &= \mathbb{E}_{\tilde{\tau}} \left[ \sum_{t=0}^{T-1} \gamma^t r_t(s_t, a_t) \mid p \right], \\ \eta_{\mathcal{D}}(\theta) &= \mathbb{E}_{p \sim \mathcal{P}} [\eta_{\mathcal{M}}(\theta, p)] \\ &= \mathbb{E}_{p \sim \mathcal{P}} \left[ \mathbb{E}_{\tilde{\tau}} \left[ \sum_{t=0}^{T-1} \gamma^t r_t(s_t, a_t) \mid p \right] \right] = \mathbb{E}_{\tau} \left[ \sum_{t=0}^{T-1} \gamma^t r_t(s_t, a_t) \right]. \end{aligned} \quad (1)$$

In (1),  $\eta_{\mathcal{M}}(\theta, p)$  is the evaluation of  $\pi_\theta$  on the model  $\mathcal{M}(p)$ , with  $\tilde{\tau}$  being trajectories generated by  $\mathcal{M}(p)$  and  $\pi_\theta$ :  $\tilde{\tau} = \{s_t, a_t, r_t\}_{t=0}^T$  where  $s_{t+1} \sim \mathcal{T}_p(s_t, a_t)$ ,  $s_0 \sim S_{0,p}$ ,  $r_t \sim \mathcal{R}_p(s_t, a_t)$ , and  $a_t \sim \pi_\theta(s_t)$ . Similarly,  $\eta_{\mathcal{D}}(\theta)$  is the evaluation of  $\pi_\theta$  over the source domain distribution. The corresponding expectation is over trajectories  $\tau$  generated by  $\mathcal{D}$  and  $\pi_\theta$ :  $\tau = \{s_t, a_t, r_t\}_{t=0}^T$ , where  $s_{t+1} \sim \mathcal{T}_{p_t}(s_t, a_t)$ ,  $p_{t+1} = p_t$ ,  $s_0 \sim S_{0,p_0}$ ,  $r_t \sim \mathcal{R}_{p_t}(s_t, a_t)$ ,  $a_t \sim \pi_\theta(s_t)$ , and  $p_0 \sim \mathcal{P}$ . With this modified notation of trajectories, batch policy optimization can be invoked for policy search.

Optimizing  $\eta_{\mathcal{D}}$  allows us to learn a policy that performs best in expectation over models in the source domain distribution. However, this does not necessarily lead to a robust policy, since there could be high variability in performance for different models in the distribution. To explicitly seek a robust policy, we use a softer version of max-min objective suggested in robust control, and optimize for the conditional value at risk (CVaR) [32]:

$$\begin{aligned} \max_{\theta, y} \quad & \int_{\mathcal{F}(\theta)} \eta_{\mathcal{M}}(\theta, p) \mathcal{P}(p) dp \\ \text{s.t.} \quad & \mathbb{P}(\eta_{\mathcal{M}}(\theta, P) \leq y) = \epsilon, \end{aligned} \quad (2)$$

where  $\mathcal{F}(\theta) = \{p \mid \eta_{\mathcal{M}}(\theta, p) \leq y\}$  is the set of parameters corresponding to models that produce the worst  $\epsilon$  percentile of returns, and provides the limit for the integral;  $\eta_{\mathcal{M}}(\theta, P)$  is the random variable of returns, which is induced by the distribution over model parameters; and  $\epsilon$  is a hyperparameter which governs the level of relaxation from max-min objective. The interpretation is that (2) maximizes the expected return for the worst  $\epsilon$ -percentile of MDPs in the source domain distribution. We adapt the previous policy gradient formulation to approximately optimize the objective in (2). The resulting algorithm, which we call EPOpt- $\epsilon$ , generalizes learning a policy

using an ensemble of source MDPs which are sampled from a source domain distribution.

---

**Algorithm 1:** EPOpt- $\epsilon$  for Robust Policy Search

---

```

1 Input:  $\psi, \theta_0, niter, N, \epsilon$ 
2 for iteration  $i = 0, 1, 2, \dots, niter$  do
3   for  $k = 1, 2, \dots, N$  do
4     sample model parameters  $p_k \sim \mathcal{P}_\psi$ 
5     sample a trajectory  $\tau_k = \{s_t, a_t, r_t, s_{t+1}\}_{t=0}^{T-1}$ 
        from  $\mathcal{M}(p_k)$  using policy  $\pi(\theta_i)$ 
6   end
7   compute  $Q_\epsilon = \epsilon$  percentile of  $\{R(\tau_k)\}_{k=1}^N$ 
8   select sub-set  $\mathbb{T} = \{\tau_k : R(\tau_k) \leq Q_\epsilon\}$ 
9   Update policy:  $\theta_{i+1} = \text{BatchPolOpt}(\theta_i, \mathbb{T})$ 
10 end

```

---

In Algorithm 1,  $R(\tau_k) \equiv \sum_{t=0}^{T-1} \gamma^t r_{t,k}$  denotes the discounted return obtained in trajectory sample  $\tau_k$ . In line 7, we compute the  $\epsilon$ -percentile value of returns from the  $N$  trajectories. In line 8, we find the subset of sampled trajectories which have returns lower than  $Q_\epsilon$ . Line 9 calls one step of an underlying batch policy optimization subroutine on the subset of trajectories from line 8. For the CVaR objective, it is important to use a good baseline for the value function. [32] show that without a baseline, the resulting policy gradient is biased and not consistent. We use a linear function as the baseline with a time varying feature vector to approximate the value function, similar to [6]. The parameters of the baseline are estimated using only the subset of trajectories with return less than  $Q_\epsilon$ . We found that this approach led to empirically good results.

For small values of  $\epsilon$ , we observed that using the sub-sampling step from the beginning led to unstable learning. Policy gradient methods adjust parameters of policy to increase probability of trajectories with high returns and reduce probability of poor trajectories. EPOpt- $\epsilon$  due to the sub-sampling step emphasizes penalizing poor trajectories more. This might constrain the initial exploration needed to find good trajectories. Thus, we initially use a setting of  $\epsilon = 1$  for few iterations before setting epsilon to the desired value. This corresponds to exploring initially to find promising trajectories and rapidly reducing probability of trajectories that do not generalize.

### B. Adapting the source domain distribution

In line with model-based Bayesian RL, we can adapt the ensemble distribution after observing trajectory data from the target domain. The Bayesian update can be written as:

$$\begin{aligned}
 \mathbb{P}(P|\tau_k) &= \frac{1}{Z} \times \mathbb{P}(\tau_k|P) \times \mathbb{P}(P) \\
 &= \frac{1}{Z} \prod_{t=0}^{T-1} \mathbb{P}(S_{t+1} = s_{t+1}^{(k)} | s_t^{(k)}, a_t^{(k)}, p) \mathbb{P}(P = p),
 \end{aligned} \tag{3}$$

where  $Z$  is the partition function (normalization) required to make the probabilities sum to 1,  $S_{t+1}$  is the random variable

representing the next state, and  $(s_t^{(k)}, a_t^{(k)}, s_{t+1}^{(k)})_{t=0}^T$  are data observed along trajectory  $\tau_k$ . We try to explain the target trajectory using the stochasticity in the state-transition function, which also models sensor errors. This provides the following expression for the likelihood:

$$\mathbb{P}(S_{t+1}|s_t, a_t, p) \equiv \mathcal{T}_p(s_t, a_t). \tag{4}$$

We follow a sampling based approach to calculate the posterior, by sampling a set of model parameters:  $p_i = [p_1, p_2, \dots, p_M]$  from a sampling distribution,  $\mathbb{P}_S(p_i)$ . Consequently, using Bayes rule and importance sampling, we have:

$$\mathbb{P}(p_i|\tau_k) \propto \mathcal{L}(\tau_k|p_i) \times \frac{\mathbb{P}_P(p_i)}{\mathbb{P}_S(p_i)}, \tag{5}$$

where  $\mathbb{P}_P(p_i)$  is the probability of drawing  $p_i$  from the prior distribution; and  $\mathcal{L}(\tau_k|p_i)$  is the likelihood of generating the observed trajectory with model parameters  $p_i$ . The weighted samples from the posterior can be used to estimate a parametric model, as we do in this paper. Alternatively, one could approximate the continuous probability distribution using discrete weighted samples like in case of particle filters. In cases where the prior has very low probability density in certain parts of the parameter space, it might be advantageous to choose a sampling distribution different from the prior. The likelihood can be factored using the Markov property as:  $\mathcal{L}(\tau_k|p_i) = \prod_t \mathbb{P}(S_{t+1} = s_{t+1}^{(k)} | s_t^{(k)}, a_t^{(k)}, p_i)$ . This simple model adaptation rule allows us to illustrate the utility of EPOpt for robust policy search, as well as its integration with model adaptation to learn policies in cases where the target model could be very different from the initially assumed distribution.

## IV. EXPERIMENTS

We evaluated the proposed EPOpt- $\epsilon$  algorithm on the 2D hopper [8] task using the MuJoCo physics simulator [35].<sup>1</sup> The batch policy optimization sub-routine is implemented using TRPO. The policy is represented with a Gaussian distribution, the mean of which is parametrized using a neural network with two hidden layers (tanh nonlinearity). Our experiments are aimed at answering the following questions: leftmargin=\*

- 1) How does the performance of standard policy search methods (like TRPO) degrade in the presence of systematic physical differences between the training and test domains, as might be the case when training in simulation and testing in the real world?
- 2) Does training on a distribution of models with EPOpt improve the performance of the policy when tested under various model discrepancies, and how much does ensemble training degrade overall performance (e.g. due to acquiring a more conservative strategy)?
- 3) How does the robustness of the policy to physical parameter discrepancies change when using the robust EPOpt- $\epsilon$  variant of our method?
- 4) Can EPOpt learn policies that are robust to unmodeled effects – that is, discrepancies in physical parameters

<sup>1</sup>Supplementary video: <https://youtu.be/w1YJ9vwaoto>

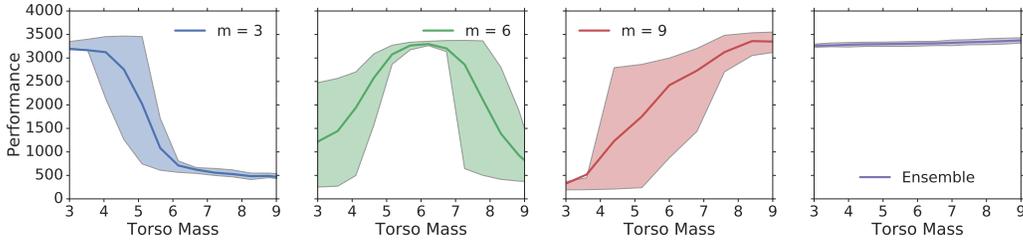


Fig. 1: Performance of trained hopper policies when testing on target domains with different torso masses. The first three plots (blue, green, and red) show the performance of policies trained with TRPO on source domains with torso mass 3, 6, and 9, respectively (denoted by  $m =$  in the legend). The rightmost plot shows the performance of EPOpt( $\epsilon = 0.1$ ) trained on a Gaussian source distribution with mean mass  $\mu = 6$  and standard deviation  $\sigma = 1.5$ . The shaded regions show the 10<sup>th</sup> and 90<sup>th</sup> percentile of the return distribution. Policies trained using traditional approaches on a single mass value are unstable for even slightly different masses, making the hopper fall over when trying to move forward. In contrast, the EPOpt policy is stable and achieves a high level of performance on the entire range of masses considered. Further, the EPOpt policy does not suffer from degradation in performance as a consequence of adopting a more robust policy.

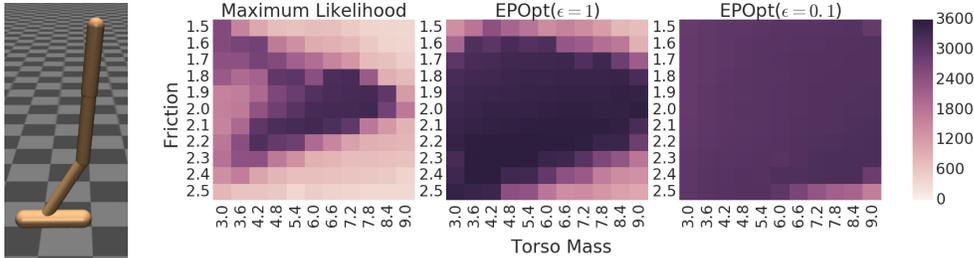


Fig. 2: On the left, is an illustration of the simulated 2D hopper task studied in this paper. On right, we depict the performance of policies for various model instances of the hopper task. The performance is depicted as a heat map for various model configurations, parameters of which are given in the x and y axis. The adversarially trained policy, EPOpt( $\epsilon = 0.1$ ), is observed to generalize to a wider range of models and is more robust. Table I.

between source and target domains that *do not* vary in the source domain ensemble?

- 5) When the initial model ensemble differs substantially from the target domain, can the ensemble be adapted efficiently, and how much data from the target domain is required for this?

In all the comparisons, *performance* refers to the average undiscounted return per trajectory or episode (we consider finite horizon episodic problems). In addition to the previously defined performance, we also use the 10<sup>th</sup> percentile of the return distribution as a proxy for the worst-case return.

#### A. Comparison to Standard Policy Search

In Figure 1, we evaluate the performance of standard TRPO and EPOpt( $\epsilon = 0.1$ ) on the hopper task, in the presence of a simple discrepancy in physics parameters between training and testing domains. The first three plots show policies that are each trained on a single torso mass in the source domain, while the last plot illustrates the performance of EPOpt, which is trained on a Gaussian mass distribution. The results show that no single torso mass value produces a policy that is successful in all target domains. However, the EPOpt policy succeeds almost uniformly for all tested mass values. Furthermore, the results show that there is almost no degradation in the performance

of EPOpt for any mass, suggesting that the EPOpt policy does not suffer substantially from adopting a more robust strategy.

#### B. Analysis of Robustness

Next, we analyze the robustness of policies trained using EPOpt on the hopper domain. For this analysis, we construct a source distribution which varies four different physical parameters: torso mass, ground friction, foot joint damping, and joint inertia (armature). This distribution is presented in Table I. Using this source distribution, we compare between three different methods: (1) standard policy search (TRPO) trained on a single model corresponding to the mean parameters in Table I; (2) EPOpt( $\epsilon = 1$ ) trained on the source distribution; (3) EPOpt( $\epsilon = 0.1$ ) – i.e. the adversarially trained policy, again trained on the previously described source distribution. The aim of the comparison is to study direct-transfer performance, similar to the robustness evaluations common in robust controller design [39]. Hence, we learn a policy using each of the methods, and then test policies on different model instances (i.e. different combinations of physical parameters) without any adaptation. The results of this comparison are summarized in Figure 2, where we present the performance of the policy for testing conditions corresponding to different torso mass and friction values, which we found to have the

most pronounced impact on performance. The results indicate that  $\text{EPOpt}(\epsilon = 0.1)$  produces highly robust policies.

TABLE I: Initial source domain distribution

Hopper	$\mu$	$\sigma$	low	high
mass	6.0	1.5	3.0	9.0
ground friction	2.0	0.25	1.5	2.5
joint damping	2.5	1.0	1.0	4.0
armature	1.0	0.25	0.5	1.5

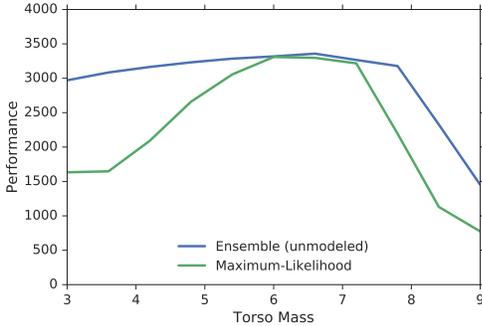


Fig. 3: Comparison between policies trained on a fixed *maximum-likelihood* model with mass (6), and an ensemble where all models have the same mass (6) and other parameters varying as described in Table I.

### C. Robustness to Unmodeled Effects

To analyze the robustness to unmodeled effects, our next experiment considers the setting where the source domain distribution is obtained by varying friction, damping, and armature as in Table I, but does not consider a distribution over torso mass. Specifically, all models in the source domain distribution have the same torso mass (value of 6), but we will evaluate the policy trained on this distribution on target domains where the torso mass is different. Figure 3 indicates that the  $\text{EPOpt}(\epsilon = 0.1)$  policy is robust to a broad range of torso masses even when its variation is not considered. However, as expected, this policy is not as robust as the case when mass is also modeled as part of the source domain distribution.

### D. Model Adaptation

The preceding experiments show that  $\text{EPOpt}$  can find robust policies, but the source distribution in these experiments was chosen to be broad enough such that the target domain is not too far from high-density regions of the distribution. However, for real-world problems, we might not have the domain knowledge to identify a good source distribution in advance. In such settings, model adaptation allows us to change the parameters of the source distribution using data gathered from the target domain. Additionally, model adaptation is helpful when the parameters of the target domain could change over time, for example due to wear and tear in a physical system. To illustrate model adaptation, we performed an experiment where the target domain was very far from the high density regions

of the initial source distribution, as depicted in Figure 4(a). In this experiment, the source distribution varies the torso mass and ground friction. We observe that progressively, the source distribution becomes a better approximation of the target domain and consequently the performance improves. In this case, since we followed a sampling based approach, we used a uniform sampling distribution, and weighted each sample with the importance weight as described in Section 3.2. Eventually, after 10 iterations, the source domain distribution is able to accurately match the target domain. Figure 4(b) depicts the learning curve, and we see that a robust policy with return more than 2500, which roughly corresponds to a situation where the hopper is able to move forward without falling down for the duration of the episode, can be discovered with just 5 trajectories from the target domain. Subsequently, the policy improves near monotonically, and  $\text{EPOpt}$  finds a good policy with just 11 episodes worth of data from the target domain. In contrast, to achieve the same level of performance on the target domain, completely model-free methods like TRPO would require more than  $2 \times 10^4$  trajectories when the neural network parameters are initialized randomly.

## V. RELATED WORK

Robust control is a branch of control theory which formally studies development of robust policies [39, 22, 17]. However, typically no distribution over source or target tasks is assumed, and a worst case analysis is performed. Most results from this field have been concentrated around linear systems or finite MDPs, which often cannot adequately model complexities of real-world tasks. The set-up of model-based Bayesian RL maintains a belief over models for decision making under uncertainty [36, 10]. In Bayesian RL, through interaction with the target domain, the uncertainty is reduced to find the correct or closest model. Application of this idea in its full general form is difficult, and requires either restrictive assumptions like finite MDPs [25], gaussian dynamics [27], or task specific innovations. Previous methods have also suggested treating uncertain model parameters as unobserved state variables in a continuous POMDP framework, and solving the POMDP to get optimal exploration-exploitation trade-off [7, 24]. While this approach is general, and allows automatic learning of epistemic actions, extending such methods to large continuous control tasks like those considered in this paper is difficult.

Risk sensitive RL methods [5, 32] have been proposed to act as a bridge between robust control and Bayesian RL. These approaches allow for using subjective model belief priors, prevent overly conservative policies, and enjoy some strong guarantees typically associated with robust control. However, their application in high dimensional continuous control tasks have not been sufficiently explored. We refer readers to [9] for a survey of related risk sensitive RL methods in the context of robustness and safety.

Standard model-based control methods typically operate by finding a maximum-likelihood estimate of the target model [18, 28, 4], followed by policy optimization. Use of model ensembles to produce robust controllers was explored recently

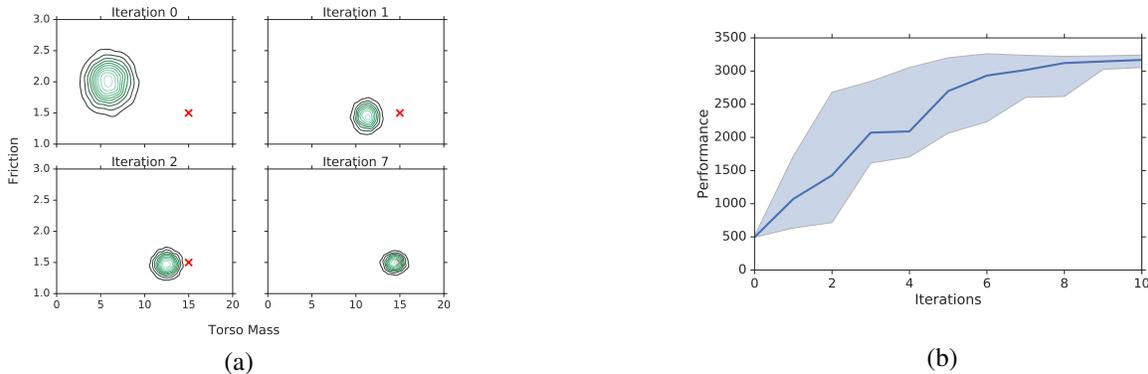


Fig. 4: (a) Visualizes the source distribution during model adaptation on the hopper task, where mass and friction coefficient are varied in the source domain. The red cross indicates the unknown parameters of the target domain. The contours in the plot indicate the distribution over models (we assume a Gaussian distribution). Lighter colors and more concentrated contour lines indicate regions of higher density. Each iteration corresponds to one round (episode) of interaction with the target domain. The high-density regions gradually move toward the true model, while maintaining probability mass over a range of parameters which can explain the behavior of target domain. Figure 4(b) presents the corresponding learning curve, where the shaded region describes the 10th and 90th percentiles of the performance distribution, and the solid line is the average performance.

in robotics. Mordatch et al. [20] use a trajectory optimization approach and an ensemble with small finite set of models; whereas we follow a sampling based direct policy search approach over a continuous distribution of uncertain parameters, and also show domain adaptation. Sampling based approaches can be applied to complex models and discrete MDPs which cannot be planned through easily. Similarly, [37] use an ensemble of models, but their goal is to optimize for average case performance as opposed to transferring to a target MDP. [37] use a hand engineered policy class whose parameters are optimized with CMA-ES. EPOpt on the other hand can optimize expressive neural network policies directly. In addition, we show model adaptation, effectiveness of the sub-sampling step ( $\epsilon < 1$  case), and robustness to unmodeled effects, all of which are important for transferring to a target MDP.

Learning of parametrized skills [3] is also concerned with finding policies for a distribution of parametrized tasks. However, this is primarily geared towards situations where task parameters are revealed during test time. Our work is motivated by situations where target task parameters (e.g. friction) are unknown. A number of methods have also been suggested to reduce sample complexity when provided with either a baseline policy [34, 13], expert demonstration [14, 2], or approximate simulator [31, 1]. These are complimentary to our work, in the sense that our policy, which has good direct-transfer performance, can be used to sample from the target domain and other off-policy methods could be explored for policy improvement.

## VI. CONCLUSIONS AND FUTURE WORK

In this paper, we presented the EPOpt- $\epsilon$  algorithm for training robust policies on ensembles of source domains. Our method provides for training of robust policies, and supports an adversarial training regime designed to provide good direct-transfer performance. We also describe how our approach can be combined with Bayesian model adaptation to adapt the

source domain ensemble to a target domain using a small amount of target domain experience. Our experimental results demonstrate that the ensemble approach provides for highly robust and generalizable policies in fairly complex simulated robotic tasks. Our experiments also demonstrate that Bayesian model adaptation can produce distributions over models that lead to better policies on the target domain than more standard maximum likelihood estimation, particularly in presence of unmodeled effects.

Although our method exhibits good generalization performance, the adaptation algorithm we use currently relies on sampling the parameter space, which is computationally intensive as the number of variable physical parameters increase. We observed that (adaptive) sampling from the prior leads to fast and reliable adaptation if the true model does not have very low probability in the prior. However, when this assumption breaks, we require a different sampling distribution which could produce samples from all regions of the parameter space. This is a general drawback of Bayesian adaptation methods. In future work, we plan to explore alternative sampling and parameterization schemes, including non-parametric distributions. In this work, we assumed access to a physics simulator that presents a compact representation of various physics parameters. However, there could be a number of factors such as geometric properties of generic objects which cannot be represented using a small number of interpretable parameters. Thus, an eventual end goal is to learn models like Bayesian neural networks, either as a replacement for a physics based model, or as an addition to explain residuals in real data which cannot be explained with the physics based model. These flexible Bayesian neural networks can be bootstrapped with physics based simulators to facilitate initial training. Such models are likely needed to train highly generalizable policies using high dimensional inputs like images or when working with hard to model objects like deformable bodies and fluids.

## REFERENCES

- [1] Pieter Abbeel, Morgan Quigley, and Andrew Y. Ng. Using inaccurate models in reinforcement learning. In *ICML*, 2006.
- [2] Brenna D. Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5):469–483, 2009.
- [3] Bruno Castro da Silva, George Konidaris, and Andrew G. Barto. Learning parameterized skills. In *ICML*, 2012.
- [4] Marc Peter Deisenroth, Gerhard Neumann, and Jan Peters. A survey on policy search for robotics. *Foundations and Trends in Robotics*, 2(12):1–142, 2013.
- [5] Erick Delage and Shie Mannor. Percentile optimization for markov decision processes with parameter uncertainty. *Operations Research*, 58(1):203–213, 2010.
- [6] Yan Duan, Xi Chen, Rein Houthoofd, John Schulman, and Pieter Abbeel. Benchmarking deep reinforcement learning for continuous control. In *ICML*, 2016.
- [7] Michael O. Duff. Design for an optimal probe. In *ICML*, 2003.
- [8] Tom Erez, Yuval Tassa, and Emanuel Todorov. Infinite-horizon model predictive control for periodic tasks with contacts. In *Proceedings of Robotics: Science and Systems*, 2011.
- [9] Javier García and Fernando Fernández. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 2015.
- [10] Mohammad Ghavamzadeh, Shie Mannor, Joelle Pineau, and Aviv Tamar. Bayesian reinforcement learning: A survey. *Foundations and Trends in Machine Learning*, 8(5-6):359–483, 2015.
- [11] Sham Kakade. A natural policy gradient. In *NIPS*, 2001.
- [12] Sham Kakade. *On the Sample Complexity of Reinforcement Learning*. PhD thesis, University College London, 2003.
- [13] Sham Kakade and John Langford. Approximately optimal approximate reinforcement learning. In *ICML*, 2002.
- [14] Sergey Levine and Vladlen Koltun. Guided policy search. In *ICML*, 2013.
- [15] Yitao Liang, Marlos C. Machado, Erik Talvitie, and Michael Bowling. State of the art control of atari games using shallow reinforcement learning. In *AAMAS*, 2016.
- [16] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra. Continuous control with deep reinforcement learning. *ArXiv e-prints*, September 2015.
- [17] Shiau Hong Lim, Huan Xu, and Shie Mannor. Reinforcement learning in robust markov decision processes. In *NIPS*. 2013.
- [18] Lennart Ljung. *System Identification*, pages 163–173. Birkhäuser Boston, Boston, MA, 1998.
- [19] Volodymyr Mnih et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, Feb 2015.
- [20] I. Mordatch, K. Lowrey, and E. Todorov. Ensemble-CIO: Full-body dynamic motion planning that transfers to physical humanoids. In *IROS*, 2015.
- [21] Igor Mordatch, Kendall Lowrey, Galen Andrew, Zoran Popovic, and Emanuel V. Todorov. Interactive control of diverse complex characters with neural networks. In *NIPS*. 2015.
- [22] Arnab Nilim and Laurent El Ghaoui. Robust control of markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798, 2005.
- [23] Xue Bin Peng, Glen Berseth, and Michiel van de Panne. Terrain-adaptive locomotion skills using deep reinforcement learning. *ACM Transactions on Graphics (Proc. SIGGRAPH 2016)*, 2016.
- [24] Josep M. Porta, Nikos A. Vlassis, Matthijs T. J. Spaan, and Pascal Poupart. Point-based value iteration for continuous pomdps. *Journal of Machine Learning Research*, 7:2329–2367, 2006.
- [25] Pascal Poupart, Nikos A. Vlassis, Jesse Hoey, and Kevin Regan. An analytic solution to discrete bayesian reinforcement learning. In *ICML*, 2006.
- [26] A. Rajeswaran, K. Lowrey, E. Todorov, and S. Kakade. Towards Generalization and Simplicity in Continuous Control. *ArXiv e-prints*, March 2017.
- [27] S. Ross, B. Chaib-draa, and J. Pineau. Bayesian reinforcement learning in continuous pomdps with application to robot navigation. In *ICRA*, 2008.
- [28] Stephane Ross and Drew Bagnell. Agnostic system identification for model-based reinforcement learning. In *ICML*, 2012.
- [29] John Schulman, Sergey Levine, Philipp Moritz, Michael Jordan, and Pieter Abbeel. Trust region policy optimization. In *ICML*, 2015.
- [30] David Silver et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, Jan 2016.
- [31] Aviv Tamar, Dotan Di Castro, and Ron Meir. Integrating a partial model into model free reinforcement learning. *Journal of Machine Learning Research*, 2012.
- [32] Aviv Tamar, Yonatan Glassner, and Shie Mannor. Optimizing the cvar via sampling. In *AAAI Conference on Artificial Intelligence*, 2015.
- [33] Matthew E. Taylor and Peter Stone. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10:1633–1685, December 2009.
- [34] Philip Thomas, Georgios Theodorou, and Mohammad Ghavamzadeh. High-confidence off-policy evaluation. In *AAAI Conference on Artificial Intelligence*. 2015.
- [35] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033, Oct 2012.
- [36] Nikos Vlassis, Mohammad Ghavamzadeh, Shie Mannor, and Pascal Poupart. *Bayesian Reinforcement Learning*, pages 359–386. Springer Berlin Heidelberg, Berlin,

Heidelberg, 2012.

- [37] Jack M. Wang, David J. Fleet, and Aaron Hertzmann. Optimizing walking controllers for uncertain inputs and environments. *ACM Trans. Graph.*, 2010.
- [38] Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3):229–256, 1992.
- [39] Kemin Zhou, John C. Doyle, and Keith Glover. *Robust and Optimal Control*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1996.