A CLOSER LOOK AT GRADIENT ESTIMATORS WITH REINFORCEMENT LEARNING AS INFERENCE

Anonymous authors
Paper under double-blind review

ABSTRACT

The concept of reinforcement learning as inference (RLAI) has led to the creation of a variety of popular algorithms in deep reinforcement learning. Unfortunately, most research in this area relies on wider algorithmic innovations not necessarily relevant to such frameworks. Additionally, many seemingly unimportant modifications made to these algorithms, actually produce inconsistencies with the original inference problem posed by RLAI. Taking a divergence minimization perspective, this work considers some of the practical merits and theoretical issues created by the choice of loss function minimized in the policy update in off-policy reinforcement learning. Our results show that while the choice of divergence rarely has a major affect on the sample efficiency of the algorithm, it can have important practical repercussions on ease of implementation, computational efficiency, and restrictions to the distribution over actions.

1 INTRODUCTION

The recent success of deep reinforcement learning (RL) across a wide range of applications is largely attributable to advances in actor-critic algorithms (20; 15; 17; 32; 42; 34; 43; 27). While versatile, many of these algorithms are difficult to implement and require careful hyper-parameter tuning (12). Additionally, the most performant implementations of these algorithms rely on heuristics to stabilize learning that are not explicitly referenced within the algorithm definition, or even informed by the relevant theory (12; 47). In some cases these implementation details become so crucial to performance that reproducing the original results becomes almost impossible (24; 11; 21).

In light of these issues, some researchers refocused on algorithms informed by probabilistic methods, most notably a popular class of algorithms derived within the 'RL as Inference' (RLAI) framework (37; 30; 29; 51; 38; 14; 2; 44; 7; 41). Under this framework, policy learning in a Markov decision process can be framed as an approximate posterior inference task, which has several key benefits over alternative approaches. First, policies trained using this approach can produce agents that transfer skills between tasks or even entire environments (13). Second, RLI couches the task of model learning (i.e. learning the dynamics of the environment) within the larger literature of generative modeling and can take advantage of existing tools and techniques (8; 2; 23; 7). Lastly, Kullback–Leibler (KL) regularization introduces the notion of "default policies", which induce different behaviors in the learned policy that are useful in solving certain environments (16; 19; 18).

Practically speaking, RLAI provides insights into how to construct various policy losses which interact more intelligently with modern computation graph software (39; 1). In framing RL as an approximate inference task, we force this consideration in deciding what divergence we wish to minimize, and how gradients of that expectation will be estimated. This is doubly important in settings where the the simulator itself can be fully or partially differentiable (5; 4). While we will not consider how gradient estimates of the forward and backwards messages can be better estimated in general, we will look at how to provide learning signal to policies in a general setting (e.g. for anonymous policy distributions given only scores of samples). We argue that closer consideration of this class of gradient estimators will allow for more composable reinforcement learning software, which in turn would allow for more stable policy learning over any model or distribution class.

This paper focuses on the following: (1) gradient estimators of the forward and reverse KL divergences with respect to the policy parameters, (2) the practical implementations of popular algorithms based on these estimators, and (3) an empirical comparison of the two resulting policy losses. The paper is

organized as follows, Section 2 reviews RL and RLAI notation and background, sections 3 and 4 derive policy gradient estimators based on divergence minimization, and section 5 considers a simple experiment to show how the performance of each objective differs under modified hyper-parameters. We conclude with a discussion on the utility of each policy loss, and consider future directions.

2 BACKGROUND

This section details three important ideas required to understand the differences between gradient estimators and target objectives in RLAI. The first is the probabilistic structure of a Markov Decision process, which will allow us to construct a corresponding inference task and model dependencies within various gradient estimators. The second is the is how actor-critic algorithms amortize many of the previous costly sampling operations required in classical inference based methods and on-policy algorithms. Lastly, we need a basic understanding of the RLAI inference task, and how it can be re-framed as divergence minimization.

2.1 MARKOV DECISION PROCESSES

We define a Markov decision process (MDP), $\mathcal{M}_\Phi(\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{T}_0, \mathcal{T}, \Pi_\Phi)$, as a random process which produces a sequence of tuples $\tau_t := \{a_t, s_t, s_{t+1}, r_t\}$, for a particular set of states $s_t \in \mathcal{S}$, actions $a_t \in \mathcal{A}$, initial state distribution $p(s_0) \in \mathcal{T}_0$, transition distribution $p(s_{t+1}|s_t, a_t) \in \mathcal{T}$, reward function $r_t : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$, and policy $\pi_\phi \in \Pi_\phi : S \to A$ parameterized by ϕ. The generative model for this finite horizon process is defined:

$$q_{\pi_\phi}(\tau) = p(s_0) \prod_{t=0}^{T} p(s_{t+1}|s_t, a_t)\pi_\phi(a_t|s_t). \tag{1}$$

Here we also denote the marginal distribution with respect to a state s, under the trajectory distribution given in equation 1, as $q_{\pi_\phi}(s)$. In this setting, RL seeks to recover the policy which maximizes the expected cumulative reward over a trajectory, $\phi^* = \arg\max_{\phi\in\Phi} \mathbb{E}_{q_{\pi_\phi}(\tau)}[\sum_{t=0}^{T} r_t(s_t, a_t)]$. More generally, we can extend this notion to infinite horizon problems through the use of a discount factor which signifies the non-zero probability of termination after every time-step. In this setting it is common to arrive at a policy through the solution to a fixed point equation (3; 46), which can be rewritten as the following optimization problem with respect to the policy parameters ϕ:

$$\phi^* = \max_{\phi\in\Phi} \mathbb{E}_{s\sim d^{\pi_\phi}, a\sim\pi_\phi}[Q^{\pi_\phi}(a, s)], \quad \text{where} \quad d^{\pi_\phi}(s) = (1-\gamma)\sum_{t=0}^{\infty} \gamma^t q_{\pi_\phi}(s_t = s), \tag{2}$$

$$\text{and} \quad Q^{\pi_\phi}(a, s) = \mathbb{E}_{p(s'|s,a)}\left[r(s, a) + \gamma \mathbb{E}_{\pi_\phi(a'|s')}\left[Q^{\pi_\phi}(a', s')\right]\right]. \tag{3}$$

Here, the distribution over states, $d^{\pi_\phi}(s)$, is referred to as the *state occupancy*. In this setting, the Q function is generally parameterized, and in most cases, learning this function is actually a much more crucial design decision then modifications made to the policy loss to ensure convergence.

2.2 SOFT ACTOR CRITIC

One popular variant of this framework is known as soft actor-critic (SAC). This algorithm is similar to most off-policy actor critic algorithms, in that it learns to estimate the expected reward ahead given the state and action (the Q-function), using the bellman error:

$$F(\theta) = \mathop{\mathbb{E}}_{(a_t, s_t, s_{t+1})\sim\mathcal{D}} \left[\left(Q_\theta^{\pi_\phi}(a_t, s_t) - \left(\hat{r}_t + \mathop{\mathbb{E}}_{a_{t+1}\sim\pi_\phi(a_{t+1}|s_{t+1})}\left[\gamma \bar{Q}_\theta^{\pi_\phi}(a_{t+1}, s_{t+1})\right]\right)\right)^2\right] \tag{4}$$

Here, D defines the set of examples gathered from the environment under a behavioral policy, or a sequence of behavioral policies. The primary distinction between SAC and other AC algorithms (34; 15) however, is that r_t is replaced with a surrogate $\hat{r}_t = r_t - \mathcal{H}_t$, where $\mathcal{H}_t$ indicates the policy entropy at state s_t. Additionally in this setting, the bar over Q indicates a target network, generally defined by the use of parameter averaging (35; 49), a model ensemble, or both (22). Using this Q function estimate, we can iteratively update our policy to match the so called "softmax-optimal"

policy (19; 20; 30). This approach defines its policy objective based on a divergence between the current policy and a normalized target Q function:

$$J_\theta(\phi) = \mathop{\mathbb{E}}_{s_t \sim \mathcal{D}} \left[\mathop{\mathbb{E}}_{a_t \sim \pi_\phi(a_t|s_t)} \left[Q_\theta^{\pi_\phi}(a_t, s_t) + \alpha \log \left(\frac{\pi_0(a_t|s_t)}{\pi_\phi(a_t|s_t)} \right) \right] \right], \tag{5}$$

Here the gradient is evaluated using the reparametrization trick (28) in cases where the distribution over actions is reparameterizable (e.g. normal, gamma ect.), and with reinforce (50) otherwise. A description of SAC is included in Algorithm 1.

Algorithm 1 Soft Actor Critic with Entropy Tuning

```
Initialize Q-functions, Q_θ1(s,a), Q_θ2(s,a) and policy weights π_φ(a_t|s_t)
Initialize target networks Q_θ̄1(s,a) , Q_θ̄2(s,a)
Initialize replay buffer D
for each iteration do
    for each environment step do
        Sample action form policy π
        store transition into replay buffer
    end for
    for each gradient step do
        Update Q-function using J_Q(θ_i) for i ∈ {1,2} θ_i ← θ_i − λ_Q ∇̂_θi J_Q(θ_i)
        Update policy weights using J_π(φ) φ ← φ − λ_π ∇̂_φ J_π(φ)
        Adjust temperature using J(α) α ← α − λ ∇̂_α J(α)
        Update target network weights using Polyak averaging θ̄_i ← τθ_i + (1 − τ)θ̄_i
    end for
end for
```

2.3 Reinforcement Learning as Inference

RL can be posed as a posterior inference problem over the set of trajectories given a notion of how well we would like the agent to do. This now breaks down into a set of latent random variables (the trajectories), and a set of observed random variables describing how well the agent is performing in the classical sense (e.g. the reward). We will refer to the observed random variables as optimality variables denoted O_t. These "optimality" variables indicate whether a given state action pair was optimal or not, and are thus assumed to be Bernoulli distributed. Somewhat intuitively, we say that a state action was optimal by indicating $O_t = 1$, while we say that a state action was sub-optimal by indicating $O_t = 0$. Each of these random variables are then assumed to be sampled from Bernoulli random variables parameterized in the following way:

$$p(O_t = 0|a_t, s_t) = 1 - \frac{1}{Z_0} \exp r(a_t, s_t), \quad p(O_t = 1|a_t, s_t) = \frac{1}{Z_0} \exp r(a_t, s_t) \tag{6}$$

We note that Z_0 is dependent on the range of values the reward can take on. In general it is assumed that the reward is strictly negative, making a simple choice for Z_0 1. With these random variables defined, we can consider the probability that a given trajectory (e.g. a sequence of state action pairs sampled from some policy interacting with the environment) is itself optimal. For finite horizon problems, this joint distribution over both latent and observed random variables is defined,

$$p(\tau, O_{1:T-1}) = p(s_1) \prod_{t=1}^{T-1} \Big[p(s_{t+1}|s_t, a_t) \pi_0(a_t|s_t) p(O_t|\tau_t) \Big]. \tag{7}$$

where $\pi_0(a_t|s_t)$ denotes the 'default' policy or prior over actions given states. For clarity, the corresponding probabilistic graphical model is given in Figure 1. However in control we don't just want a metric on how likely a given trajectory is to be optimal, we want to produce a policy which is itself induces trajectories that are optimal: $\pi^*(a_t|s_t) = \pi(a_t|s_t, o_t, ..., o_{T-1})$. Unfortunately, arriving at an exact $\pi^*(a_t|s_t)$ is intractable in all but the simplest cases. Whats more, both our graphical model and as a resulting joint distribution are defined in terms of entire trajectories. This means that not only must we perform a posterior inference task (e.g. determine what the distribution over

states and actions given optimality), but we must also find a way recover the marginal distribution over actions given state and optimality for any state which we might see in practice. We will address the first issue (intractability of the posterior inference problem) by using variational inference to produce an approximate posterior efficiently. We will then show that through manipulation of the expected gradient and structure of the graphical model, we can define an objective who's optimum characterizes the same policy as is defined by $\pi(a_t|s_t, o_t, ..., o_{T-1})$.

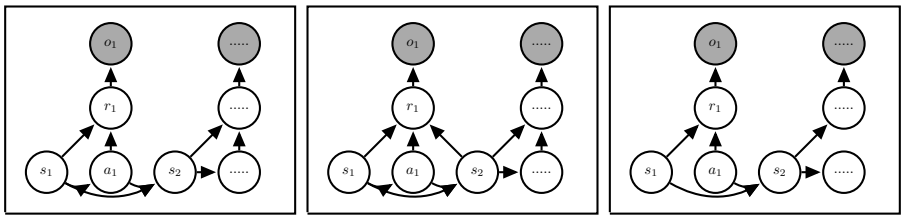

Figure 1: Three potential probabilistic graphical models associated with RLAI. In the graphs, shaded nodes indicate observed random variables, while unobserved or *latent* random variables are transparent. Note that on the left, as is usually the case the reward is not dependent on the resulting transition. The center graphical model gives an example of this, and is often not used as it tends to exacerbate over-optimistic behavior in the policy. The case on the right is given in the "maximum entropy" setting where the default policy is assumed to be uniform and independent of the state.

3 Policy Gradient Estimators Via The Reverse KL Divergence

This section contains proofs and derivations relevant to the reverse KL applied to RLAI. More specifically, this section gives information on the un-marginalized objective, the marginalized objective, the REINFORCE trick (50), and the reparameterization trick (26). This section will rely on many of the usual tools found in variational inference, and makes heavy use of the REINFORCE trick (50), and zero expectation score function trick (45).

3.0.1 The Un-marginalized Reverse KL Objective

For the following we assume that $\tau_t := (a_t, s_t, s_{t+1}, r_t)$, and $(s_{0:T}, a_{0:T-1}, r_{0:T-1}) := \tau$. We begin by looking at the objective derived by integrating over all time-steps simultaneously and differentiating. To derive this objective we start with the functional form of the reverse KL divergence,

$$\min_{\phi} \quad D_{KL}(q_\phi(\tau) \mid\mid p(\tau|O)) = \min_{\phi} \left[\int_\tau q_\phi(\tau) \log \left(\frac{q_\phi(\tau)}{p(\tau, O)} \right) d\tau \right]$$

Next we can look at the log term and note, that it can be broken up into each time step, and additionally that the dynamics between the the joint and approximate distribution cancel out. Below we assume that anywhere O is written, we have $O_{0:T-1} = 1$ (i.e. the posterior is over optimal policies),

$$\log \left(\frac{q_\phi(\tau)}{p(\tau, O)} \right) = - \sum_{t=0}^{T-1} \left(r(a_t, s_t) - \frac{\log(\pi_0(a_t|s_t))}{\log(\pi_\phi(a_t|s_t))} \right).$$

With this expression in mind, we can simplify the integral to exclude terms which are constant with respect to the policy parameters. This allows to convert a difficult to compute objective which included the optimal policy, into an objective we can estimate through Monte-Carlo sampling and interaction with the environment:

$$\min_{\phi} \left[\int_\tau q_\phi(\tau) \log \left(\frac{q_\phi(\tau)}{p(\tau, O)} \right) d\tau \right] = \min_{\phi} - \mathop{\mathbb{E}}_{\tau \sim q_\phi} \left[\sum_{t=0}^{T-1} r(s_t, a_t) - \log \frac{\pi_\phi(a_t|s_t)}{\pi_0(a_t|s_t)} \right].$$

3.0.2 The Marginalized Reverse KL gradient

Unfortunately, the estimator derived above can prove to have high variance in practice. However using this gradient estimator we can infer an objective whose minimum is the same as the one above

up to a constant factor. In a similar fashion to the policy gradient theorem, we can show the following:

$$\nabla_\phi \mathop{\mathbb{E}}_{q(\tau)} \left[\sum_{t=0}^{T-1} r(s_t, a_t) - \log \frac{\pi_\phi(a_t|s_t)}{\pi_0(a_t|s_t)} \right] = \sum_{t=0}^{T-1} \mathop{\mathbb{E}}_{q_\phi(\tau_{t:T-1})} \left[\hat{\nabla}_\phi \left(\sum_{t'=t}^{T} r(s_{t'}, a_{t'}) - \log \frac{\pi_\phi(a_{t'}|s_{t'})}{\pi_0(a_{t'}|s_{t'})}) \right) \right],$$

where $\hat{\nabla}_\phi = \nabla_\phi \log \pi_\phi(s_t, a_t)$. We start by passing the gradient into the expectation under the assumption that the gradient can be passed into the integral, which should hold so long as the expectation is bounded.

$$\begin{aligned}
\nabla_\phi J_{q,p}(\phi) &= \int \nabla_\phi \left[\left(\sum_{t=0}^{T-1} r(s_t, a_t) - \log \frac{\pi_\phi(a_t|s_t)}{\pi_0(a_t|s_t)} \right) q_\phi(\tau) \right] d\tau \\
&= \int \left(\sum_{t=0}^{T-1} \nabla_\phi \log \pi_\phi(a_t|s_t) \right) \left(\sum_{t'=0}^{T-1} r(s_{t'}, a_{t'}) - \log \frac{\pi_\phi(a_{t'}|s_{t'})}{\pi_0(a_{t'}|s_{t'})} - 1 \right) q_\phi(\tau) d\tau
\end{aligned}$$

Note that s_t and s'_t correspond to the *same* random variable (as do a_t and a'_t for all t. The primes are introduced only to index the double sum below correctly. Here we use the score function trick to transform the integral of a gradient into the expectation over a gradient. We then combine terms and simplify, noting that the gradient of the dynamics is zero. We then convolve these sums and further simplify using d separation of the probabilistic graphical model.

$$\begin{aligned}
J_{q,p}(\phi) &= \int \sum_{t=0}^{T-1} \sum_{t'=0}^{T-1} \nabla_\phi \log \pi_\phi(a_t|s_t) \left(r(s_{t'}, a_{t'}) - \log \frac{\pi_\phi(a_{t'}|s_{t'})}{\pi_0(a_{t'}|s_{t'})} - 1 \right) q_\phi(\tau) d\tau, \\
&= \sum_{t=0}^{T-1} \sum_{t'=t}^{T-1} \int \nabla_\phi \log \pi_\phi(a_t|s_t) \left(r(s_{t'}, a_{t'}) - \log \frac{\pi_\phi(a_{t'}|s_{t'})}{\pi_0(a_{t'}|s_{t'})} \right) q_\phi(\tau) d\tau.
\end{aligned}$$

Where the last line follows from the Markov property of the graphical model which gives $t' < t \rightarrow r(s_{t'}, a_{t'}) - \log \pi_\phi(a_{t'}|s_{t'}) + \log \pi_0(a_{t'}|s_{t'})$ is independent of $\nabla_\phi \log \pi_\phi(a_t|s_t)$. This means that we can decompose the expectation such that for all $t' < t$,

$$\int \nabla_\phi \log \pi_\phi(a_t|s_t) f(a_{t'}, s_{t'}) q_\phi(\tau) d\tau = \int f(a_{t'}, s_{t'}) \, (0) \, \frac{q_\phi(\tau)}{\pi_\phi(a_t|s_t)} d\tau \setminus a_t = 0$$

Therefore the inner expectation is a constant for all t' less then t, and the score function estimator is zero in expectation at these terms in the series. This sum of (entropy regularized) rewards ahead is often referred to as the Q function or the advantage (if we include a baseline or value function which integrates over all actions). This term can be defined more explicitly as,

$$Q(s_t, a_t) = \mathop{\mathbb{E}}_{q_\phi(\tau_{t+1:T})|\tau_t = \{a_t, s_t\}} \left[\sum_{t'=t}^{T} r(s_{t'}, a_{t'}) - \log \frac{\pi_\phi(a_{t'}|s_{t'})}{\pi_0(a_{t'}|s_{t'})} \right].$$

Using this estimate of the expected rewards ahead, we can re-scale the log probability of certain actions and give them more weight in the gradient estimator. Crucially, the variance of this gradient estimator is strictly lower, while maintaining the correct target distribution of interest. The full marginalized objective is given as the following:

$$\nabla_\phi J_{q,p}(\phi) = \sum_{t=0}^{T-1} \mathbb{E}_{q_\phi(a_t, s_t)} \left[\nabla_\phi \log \pi_\phi(a_t|s_t) Q(s_t, a_t) \right].$$

4 Reweighted Policy Gradient Estimators Via the Forward KL

This section details derivations for how to produce an approximate optimal policy under the forward KL, and like the previous section we will produce a lower variance gradient estimator which targets the original distribution of interest. For this KL, we take advantage of tools from self normalized importance weighting to produce our estimator, and briefly discuss some of the issues present.

4.0.1 UN-MARGINALIZED OBJECTIVE DERIVATION FOR THE FORWARD KL

This objective, represents the expected KL divergence between a target distribution, and an approximate distribution under samples from the target distribution given by $KL(p(\tau|O)||q_\phi(\tau))$. Again, our latent variables are given as the associated state-action pairs, while the observed variables represents the optimality distribution. In order to evaluate the expression above, we apply a standard importance weighting scheme, beginning with the removal constant terms within the expectation.

$$KL(p(\tau|O)||q_\phi(\tau)) = \int p(\tau|O)\log\frac{p(\tau|O)}{q_\phi(\tau)}d\tau = -\int p(\tau|O)\log q_\phi(\tau)d\tau + \int p(\tau|O)\log p(\tau|O)d\tau$$
$$= \int \frac{1}{p(O)}\log q_\phi(\tau)\frac{p(\tau,O)}{q_\phi(\tau)}q_\phi(\tau)d\tau + c = \mathop{\mathbb{E}}_{\tau\sim q_\phi}\left[\frac{1}{Z}\log q_\phi(\tau)\frac{p(\tau,O)}{q_\phi(\tau)}\right] + c$$

4.0.2 UN-MARGINALIZED GRADIENT OF THE FORWARD KL OBJECTIVE

Crucially, because the original problem is over a distribution not dependent on ϕ, we can directly pass the gradient into the expectation without the REINFORCE trick, or the reparameterization trick. This gives the following expression for the gradient:

$$\nabla_\phi KL(p(\tau|O=1)||q_\phi(\tau)) = \mathop{\mathbb{E}}_{\tau\sim q_\phi}\left[\frac{p(\tau,O)}{q_\phi(\tau)}\frac{\nabla_\phi \log q_\phi(\tau)}{Z}\right]$$

This actually gives us a simple algorithm that can evaluate the gradient, and thereby minimize the expected KL between our two distributions following samples generated from our policy interacting with the simulator. Based upon this approach, we can apply self normalized importance weighting to avoid explicitly computing the constant Z. This is done in the following way:

$$\nabla_\phi KL(p(\tau|O)||q_\phi(\tau)) \approx \sum_{i=1}^{m}\left[w_i \nabla_\phi \log q_\phi(\tau_i)\right] = \frac{1}{Z}\sum_{i=1}^{m}\left[\frac{p(\tau_i,O_i)}{q_\phi(\tau_i)}\nabla_\phi \log q_\phi(\tau_i)\right].$$

Where for optimal trajectories ($O = O_{t=0:T-1}$), the un-normalized weights are defined as,

$$w_i = \frac{1}{Z}\prod_{t=0}^{T-1}\frac{\frac{1}{Z_0}\exp r_t(a_t,s_t)\pi_0(a_t|s_t)}{\pi_\phi(a_t|s_t)} = \frac{1}{\hat{Z}}\exp\left[\sum_{t=0}^{T-1} r(a_t^i,s_t^i) - \log\left(\frac{\pi_\phi(a_t^i|s_t^i)}{\pi_0(a_t^i|s_t^i)}\right)\right]$$

Where the first line follows from cancellation of the dynamics, and the second follows from the fact that we are only considering $O_t = 1$. In order to avoid computation of Z explicitly, we take advantage of the consistent but biased self normalized importance weighting estimator. This means that the weights w_i are replaced with the following:

$$\hat{w_i} = \frac{w_i}{\sum_{j=1}^{m} w_j} = \frac{\exp\left[\sum_{t=0}^{T-1} r(a_t^i,s_t^i) - \log\left(\frac{\pi_\phi(a_t^i|s_t^i)}{\pi_0(a_t^i|s_t^i)}\right)\right]}{\sum_{j=1}^{m}\exp\left[\sum_{t=0}^{T-1} r(a_t^j,s_t^j) - \log\left(\frac{\pi_\phi(a_t^j|s_t^j)}{\pi_0(a_t^j|s_t^j)}\right)\right]}. \tag{8}$$

In this case we know now have a more biased estimator of the gradient, but also one that is low variance when compared to its un-marginalized RKL counterpart (29),

$$\nabla_\phi KL(p(\tau|O)||q_\phi(\tau)) \approx \sum_{i=1}^{n}\frac{\exp\left[\sum_{t=0}^{T-1} r(a_t^i,s_t^i) - \log\left(\frac{\pi_\phi(a_t^i|s_t^i)}{\pi_0(a_t^i|s_t^i)}\right)\right]}{\sum_{j=1}^{n}\exp\left[\sum_{t=0}^{T-1} r(a_t^j,s_t^j) - \log\left(\frac{\pi_\phi(a_t^j|s_t^j)}{\pi_0(a_t^j|s_t^j)}\right)\right]}\nabla_\phi \log q_\phi(\tau_i) \tag{9}$$

4.0.3 MARGINALIZATION PROOF FOR FORWARD KL

In order to show that the objective can be marginalized with respect to actions given states and optimality, we start with the original definition of the objectives gradient and simplify until we have produced a lower variance estimator:

$$\nabla_\phi J_{p,q}(\phi) = \sum_{t=0}^{T-1}\int_\tau p(a_0,s_0,...s_T|o_t,...,o_{T-1},s_t)p(s_t|o_0,...,o_{T-1})\nabla_\phi \log \pi_\phi(a_t|s_t)d\tau.$$

Where the last line uses d separation of the probabilistic graphical model to ignore the optimality variables prior to the state being conditioned on. Next if we look at just one element of this sum at index t,

$$
\begin{aligned}
&= \int_{s_t} ... \int_{s_T} \int_{a_t} ... \int_{a_{T-1}} p(s_t|O_0...O_{T-1}) \nabla_\phi \log \pi_\phi(a_t|s_t) \\
&\quad ... \int_{s_0} ... \int_{s_{t-1}} \int_{a_0} ... \int_{a_{t-1}} p(a_0, s_0, ...s_T|o_t, ..., o_{t-1}, s_t) d\tau \\
&= \int_{s_t} ... \int_{s_T} \int_{a_t} ... \int_{a_{T-1}} p(s_t|O_0...O_{T-1}) \nabla_\phi \log \pi_\phi(a_t|s_t) p(s_{t+1}, a_t...s_T|o_t, ..., o_{t-1}, s_t) d\tau
\end{aligned}
$$

Here we integrate out all density that is un-affected by the argument within the expectation. We start by approximating the first term with importance weighting:

$$
p(s_t = \hat{s}_t|O_0...O_{T-1}) = \int_\tau \frac{1}{Z} \frac{p(a_0, s_0, ...s_T, o_0, ..., o_{T-1}, \hat{s}_t)}{q_\phi(a_0, s_0, ...s_T, o_0, ..., o_{T-1}, \hat{s}_t)} d\tau_{s_t \notin} \approx \sum_{i=1}^{n} \frac{y_i}{\sum_{i=1}^{n} y_j}
$$

where we define the weights here as:

$$
y_i = \begin{cases} \exp\left[\sum_{t=0}^{T-1} r(a_t, s_t) - \log \pi_\phi(a_t|s_t)\right], & \text{if } s_t^i = \hat{s}_t, \\ 0, & \text{else.} \end{cases}
$$

This definition allows us to define gradient in the following way:

$$
\begin{aligned}
\nabla_\phi J_{p,q}(\phi) &= \int_{s_t} ... \int_{s_T} \int_{a_t} ... \int_{a_{T-1}} p(s_t|O_0...O_{T-1}) \nabla_\phi \log \pi_\phi(a_t|s_t) p(s_{t+1}, a_t...s_T|o_t, ..., o_{t-1}, s_t) d\tau \\
&= \mathop{\mathbb{E}}_{s_t|O_{0:T-1} \sim q_\phi} \left[\frac{p(s_t|O_0...O_{T-1})}{q_\phi(s_t)} ... \right. \\
&\quad \left. ... \mathop{\mathbb{E}}_{a_{t:T-1}, s_{t+1:T}|O_{t:T-1}, s_t \sim q_\phi} \left[\nabla_\phi \log \pi_\phi(a_t|s_t) \frac{p(s_{t+1}, a_t...s_T|o_t, ..., o_{t-1}, s_t)}{q_\phi(s_{t+1}, a_t...s_T|o_t, ..., o_{t-1}, s_t)} \right] \right]
\end{aligned}
$$

We can now define our importance weights as:

$$
\begin{aligned}
y_k^t &= \frac{1}{Z} \frac{p(s_t^k|O_0...O_{T-1})}{q_\phi(s_t^k)} = \exp\left[\sum_{t=0}^{T-1} r(a_t, s_t) - \log\left(\frac{\pi_\phi(a_t^i|s_t^i)}{\pi_0(a_t|s_t)}\right)\right] \\
w_{i,k}^t &= \frac{1}{Z} \frac{p(s_{t+1}^i, a_t^i...s_T^i|o_t, ..., o_{T-1}, s_t^k)}{q_\phi(s_{t+1}^i, a_t^i...s_T^i|o_t, ..., o_{T-1}, s_t^k)} \\
&= \exp\left[r(a_t^i, s_t^k) - \log\left(\frac{\pi_\phi(a_{t'}^i|s_{t'}^i)}{\pi_0(a_{t'}^i|s_{t'}^i)}\right) + \sum_{t'=t+1}^{T-1} r(a_t^i, s_t^i) - \log\left(\frac{\pi_\phi(a_{t'}^i|s_{t'}^i)}{\pi_0(a_{t'}^i|s_{t'}^i)}\right)\right]
\end{aligned}
$$

Because we have parameterized an entropy regularized Q function, the second expression can be written more simply in the following way:

$$
w_{i,k}^t = \exp\left[r(a_t^i, s_t^k) - \log\left(\frac{\pi_\phi(a_t^i|s_t^k)}{\pi_0(a_t^i|s_t^k)}\right) + \sum_{t'=t+1}^{T-1} r(a_t^i, s_t^i) - \log\left(\frac{\pi_\phi(a_{t'}^i|s_{t'}^i)}{\pi_0(a_{t'}^i|s_{t'}^i)}\right)\right] = \exp\left[Q(a_t^i, s_t^k)\right]
$$

This leaves us with a simple gradient estimator which we can use to directly update the policy to approximate our target posterior distribution:

$$
\nabla_\phi J_{p,q}(\phi) = \sum_{t=0}^{T-1} \sum_{k=1}^{m} \frac{y_k^t}{\sum_{j=1}^{n} y_j^t} \left[\sum_{i=1}^{n} \left(\frac{w_{i,k}^t}{\sum_{j=1}^{n} w_{j,k}^t}\right) \nabla_\phi \log \pi_\phi(a_t^i|s_t^k)\right] \tag{10}
$$

$$
\text{where} \quad y_k^t = \left[\sum_{t=0}^{T-1} r(a_t, s_t) - \log\left(\frac{\pi_\phi(a_t^i|s_t^i)}{\pi_0(a_t|s_t)}\right)\right] \quad \text{and} \quad w_{i,k}^t = \exp\left[Q(a_t^i, s_t^k)\right]
$$

It should be noted here that the estimate of the marginal state distribution under the optimal policy is in fact even harder to compute that the optimal policy, as it requires particle-smoothing. Additionally, it is not clear that the policy should be weighted such that it only acts optimally in areas which are highly likely to be optimal. There most practical algorithms that use this policy objective tend to ignore the outer expectation, and simply apply updates to all states seen thus far during training uniformly. For this reason, we do not include this term in the policy objective used below in our experiments. We believe that this is the reason that most instances of this objective in the deep RL literature are referred to as optimistic, while if the state-occupancy term is ignored, you recover an objective which with sufficient functional approximation power, and enough particles will recover the policy which maximizes return.

5 Experiments

To display the relative merits of both of these objectives, we include a small set of experiments where a native implementation of soft actor critic was taken, and modified to include the forward KL objective (given below as FKL). We begin with an illustration of an instance where you can find improvement by swapping the more common RKL for the FKL objective.

In Figure 2, we plot the results of training on the Humanoid benchmark from MuJoCo using RKL and FKL. Here, we achieve state-of-the-art sample efficiency, demonstrating approximately 60% increase in sample efficiency. Humanoid is the most challenging of the five tasks considered from the MuJoCo suite. Whats more, because the FKL is distribution agnostic, meaning that it can be combined with any policy distribution without needing to be modified, the additional coding overhead was negligible. Whereas for the RKL, had we wished to implement it for say a gamma or mixture model, we would have had to modify our objective twice.

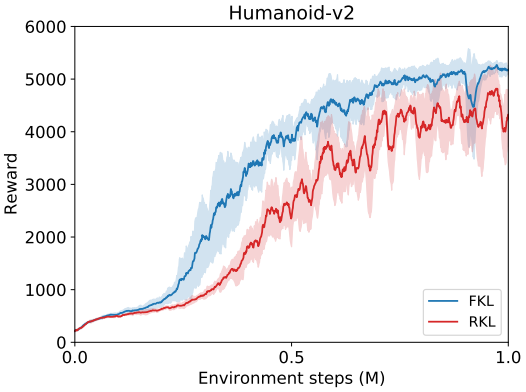

Figure 2: Comparison of RKL and FKL on the MuJoCo Humanoid benchmark. Error bars denote one standard deviation. As shown, the reward garnered throughout optimization is also lower variance trained with FKL than RKL.

For continuous control, we additionally show that the FKL can also improve the overall algorithms robustness to changes in hyper-parameters. All experiments are run with at 10 replicates and plotted with one-half standard deviation and smoothed using a 0.9 weighted EMA unless otherwise stated. As can be seen in Figure 3, the forward KL can maintain policy performance, even when memory is severely limited.

6 Related Work

There is a growing literature on various methods that incorporate the RLI framework discussed within this paper (2; 19; 30; 51; 52; 25; 37; 48). In many cases these works use RLI to give theoretical basis for the empirical effectiveness of well established implementation tricks such as entropy bonuses (19; 30; 51). Very few works however, have considered the comparison between the forward and reverse KL as well as the relative merits of each in the context of reinforcement learning objectives (37; 33). Even fewer produce scalable and sample efficient implementations of both objectives, making a proper comparison difficult.

Algorithmically, the closest class of algorithms to our work in the classical literature comes from policy search methods (37; 36; 31; 9; 40; 10). Of the subset of these that have been made scalable, closely related algorithms to our work come from (2; 44; 7). Unlike (2), we do not require an alternating update and directly learn the policy via the forward KL. Additionally while we do not provide monotonic improvement guarantees which require additional assumptions and a trust-region update, we do provide a full proof of marginalization of the forward KL as it applies to RLI in a finite

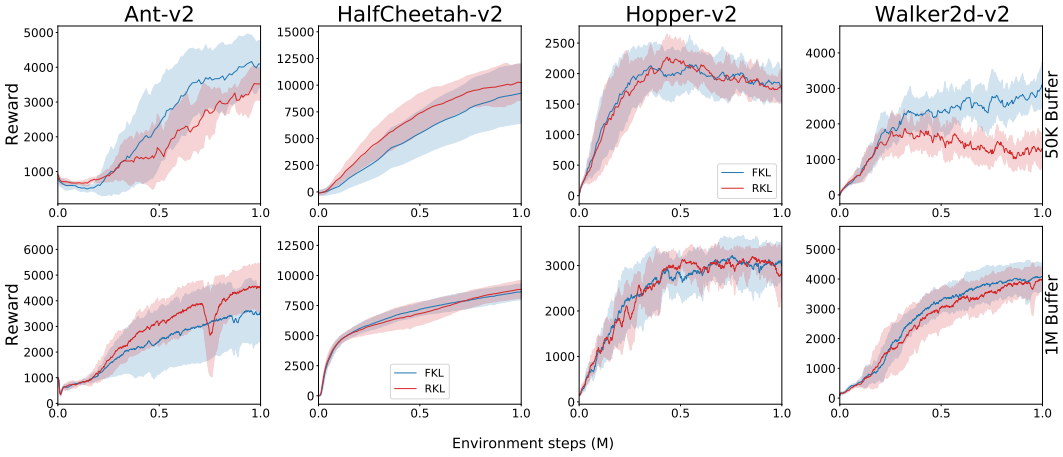

Figure 3: First row: comparison of RKL and FKL on MuJoCo benchmark tasks using a base implementation of SAC *without* hyper-parameter tuning for MuJoCo. The notable difference in hyper-parameters between these settings is the buffer sizes of 50K versus 1M for the first and second rows, respectively. We find under these conditions, FKL is always competitive or significantly better than the baseline with the exception of the Ant benchmark using 1M replay buffer size. In all figures, error bars denote one-quarter standard deviation.

horizon setting. A similar proof for a marginalization can be found in (41), however this assumes access to the steady state distribution of the Markov chain as well as a model of the environment. We do note that the objective for the target policy given in (44) does come very close to our own, however it requires additional parameterization of a value function, as well as a more complicated alliterating update for the temperature coefficient α, and a different inference formulation which culminates in an advantage based objective. Most importantly, while these works represent end-to-end algorithms, our modifications that can be applied a wide variety of existing frameworks.

Lastly with respect to our discussion of informative priors, this work was originally derived from work by (6). The key distinction being that in RLI, our observations are static, and thus our inference task slightly different. There is also a connection to the work done by (16) which in a similar way to our own algorithm, simultaneously update a prior over the action distribution within the KL regularized RL setting. This work however, focuses on inducing re-usable behaviors in hierarchical models to accelerate learning, where the method that we discuss can be applied in a wide range of algorithms and scenarios to induce whatever type of learned behavior is desired.

7 Discussion

This paper describes in detail the differences in objective which occur both mathematically, and practically under a divergence minimization perspective of RLAI. Crucially, we show that there are very real advantages to considering other objectives besides the classic policy gradient methods which are most common. While these alternatives like the FKL are generally not going to drastically change performance, their ease of implementation, robustness to hyperparameters, and their distribution agnostic nature make them a very desirable alternative to simple scalable RL test-beds. In the future we would like to see how these objectives, as well as their symmetric counterparts function on other benchmarks using other types of policy distributions.

References

[1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek

Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.

[2] Abbas Abdolmaleki, Jost Tobias Springenberg, Yuval Tassa, Remi Munos, Nicolas Heess, and Martin Riedmiller. Maximum a posteriori policy optimisation. *arXiv preprint arXiv:1806.06920*, 2018.

[3] Dimitri P Bertsekas. *Reinforcement learning and optimal control*. Athena Scientific, 2019.

[4] Jeff Bezanson, Stefan Karpinski, Viral B Shah, and Alan Edelman. Julia: A fast dynamic language for technical computing. *arXiv preprint arXiv:1209.5145*, 2012.

[5] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018.

[6] Tamara Broderick, Nicholas Boyd, Andre Wibisono, Ashia C Wilson, and Michael I Jordan. Streaming variational bayes. In *Advances in neural information processing systems*, pages 1727–1735, 2013.

[7] Yinlam Chow, Brandon Cui, MoonKyung Ryu, and Mohammad Ghavamzadeh. Variational model-based policy optimization. *arXiv preprint arXiv:2006.05443*.

[8] Ignasi Clavera, Jonas Rothfuss, John Schulman, Yasuhiro Fujita, Tamim Asfour, and Pieter Abbeel. Model-based reinforcement learning via meta-policy optimization, 2018.

[9] Adriá Colomé and Carme Torras. Dual reps: A generalization of relative entropy policy search exploiting bad experiences. *IEEE Transactions on Robotics*, 33(4):978–985, 2017.

[10] Marc Peter Deisenroth, Gerhard Neumann, and Jan Peters. *A survey on policy search for robotics*. now publishers, 2013.

[11] Yan Duan, Xi Chen, Rein Houthooft, John Schulman, and Pieter Abbeel. Benchmarking deep reinforcement learning for continuous control, 2016.

[12] Logan Engstrom, Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Firdaus Janoos, Larry Rudolph, and Aleksander Madry. Implementation matters in deep rl: A case study on ppo and trpo. In *International Conference on Learning Representations*, 2019.

[13] Matthew Fellows, Anuj Mahajan, Tim G. J. Rudner, and Shimon Whiteson. Virel: A variational inference framework for reinforcement learning, 2018.

[14] Matthew Fellows, Anuj Mahajan, Tim GJ Rudner, and Shimon Whiteson. Virel: A variational inference framework for reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 7120–7134, 2019.

[15] Scott Fujimoto, Herke Van Hoof, and David Meger. Addressing function approximation error in actor-critic methods. *arXiv preprint arXiv:1802.09477*, 2018.

[16] Alexandre Galashov, Siddhant M Jayakumar, Leonard Hasenclever, Dhruva Tirumala, Jonathan Schwarz, Guillaume Desjardins, Wojciech M Czarnecki, Yee Whye Teh, Razvan Pascanu, and Nicolas Heess. Information asymmetry in kl-regularized rl. *arXiv preprint arXiv:1905.01240*, 2019.

[17] Shixiang Gu, Timothy Lillicrap, Zoubin Ghahramani, Richard E Turner, and Sergey Levine. Q-prop: Sample-efficient policy gradient with an off-policy critic. *arXiv preprint arXiv:1611.02247*, 2016.

[18] Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. Reinforcement learning with deep energy-based policies, 2017.

[19] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.

[20] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, and Sergey Levine. Soft actor-critic algorithms and applications, 2018.

[21] Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger. Deep reinforcement learning that matters, 2017.

[22] Matteo Hessel, Joseph Modayil, Hado Van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan Horgan, Bilal Piot, Mohammad Azar, and David Silver. Rainbow: Combining improvements in deep reinforcement learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

[23] Maximilian Igl, Luisa Zintgraf, Tuan Anh Le, Frank Wood, and Shimon Whiteson. Deep variational reinforcement learning for pomdps. *arXiv preprint arXiv:1806.02426*, 2018.

[24] Riashat Islam, Peter Henderson, Maziar Gomrokchi, and Doina Precup. Reproducibility of benchmarked deep reinforcement learning tasks for continuous control, 2017.

[25] Hilbert J Kappen, Vicenç Gómez, and Manfred Opper. Optimal control as a graphical model inference problem. *Machine learning*, 87(2):159–182, 2012.

[26] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[27] Thanard Kurutach, Ignasi Clavera, Yan Duan, Aviv Tamar, and Pieter Abbeel. Model-ensemble trust-region policy optimization. *arXiv preprint arXiv:1802.10592*, 2018.

[28] Tuan Anh Le, Maximilian Igl, Tom Rainforth, Tom Jin, and Frank Wood. Auto-encoding sequential monte carlo. *arXiv preprint arXiv:1705.10306*, 2017.

[29] Tuan Anh Le, Adam R. Kosiorek, N. Siddharth, Yee Whye Teh, and Frank Wood. Revisiting reweighted wake-sleep for models with stochastic control flow, 2018.

[30] Sergey Levine. Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909*, 2018.

[31] Sergey Levine and Vladlen Koltun. Guided policy search. In *International Conference on Machine Learning*, pages 1–9, 2013.

[32] 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.

[33] Xubo Lyu and Mo Chen. Ttr-based rewards for reinforcement learning with implicit model priors. *arXiv preprint arXiv:1903.09762*, 2019.

[34] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937, 2016.

[35] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.

[36] William H Montgomery and Sergey Levine. Guided policy search via approximate mirror descent. In *Advances in Neural Information Processing Systems*, pages 4008–4016, 2016.

[37] Gerhard Neumann et al. Variational inference for policy search in changing situations. In *Proceedings of the 28th International Conference on Machine Learning, ICML 2011*, pages 817–824, 2011.

[38] Brendan O'Donoghue, Ian Osband, and Catalin Ionescu. Making sense of reinforcement learning and probabilistic inference. *arXiv preprint arXiv:2001.00805*, 2020.

[39] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.

[40] Jan Peters, Katharina Mulling, and Yasemin Altun. Relative entropy policy search. In *Twenty-Fourth AAAI Conference on Artificial Intelligence*, 2010.

[41] Alexandre Piché, Valentin Thomas, Cyril Ibrahim, Yoshua Bengio, and Chris Pal. Probabilistic planning with sequential monte carlo methods. In *International Conference on Learning Representations*, 2019.

[42] 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.

[43] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017.

[44] H Francis Song, Abbas Abdolmaleki, Jost Tobias Springenberg, Aidan Clark, Hubert Soyer, Jack W Rae, Seb Noury, Arun Ahuja, Siqi Liu, Dhruva Tirumala, et al. V-mpo: On-policy maximum a posteriori policy optimization for discrete and continuous control. *arXiv preprint arXiv:1909.12238*, 2019.

[45] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press, 2011.

[46] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. A Bradford Book, Cambridge, MA, USA, 2018.

[47] 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.

[48] Jan-Willem Van De Meent, Brooks Paige, David Tolpin, and Frank Wood. Black-box policy search with probabilistic programs. 2016.

[49] Ziyu Wang, Tom Schaul, Matteo Hessel, Hado Hasselt, Marc Lanctot, and Nando Freitas. Dueling network architectures for deep reinforcement learning. In *International conference on machine learning*, pages 1995–2003, 2016.

[50] Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3):229–256, May 1992.

[51] Brian D. Ziebart. *Modeling Purposeful Adaptive Behavior with the Principle of Maximum Causal Entropy*. PhD thesis, Carnegie Mellon University, USA, 2010.

[52] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, pages 1433–1438. Chicago, IL, USA, 2008.