
Skew-Fit: State-Covering Self-Supervised Reinforcement Learning

Vitchyr H. Pong^{*1} Murtaza Dalal^{*1} Steven Lin^{*1} Ashvin Nair¹ Shikhar Bahl¹ Sergey Levine¹

Abstract

Autonomous agents that must exhibit flexible and broad capabilities will need to be equipped with large repertoires of skills. Defining each skill with a manually-designed reward function limits this repertoire and imposes a manual engineering burden. Self-supervised agents that set their own goals can automate this process, but designing appropriate goal setting objectives can be difficult, and often involves heuristic design decisions. In this paper, we propose a formal exploration objective for goal-reaching policies that maximizes state coverage. We show that this objective is equivalent to maximizing goal reaching performance together with the entropy of the goal distribution, where goals correspond to full state observations. To instantiate this principle, we present an algorithm called Skew-Fit for learning a maximum-entropy goal distributions and prove that, under regularity conditions, Skew-Fit converges to a uniform distribution over the set of valid states, even when we do not know this set beforehand. Our experiments show that combining Skew-Fit for learning goal distributions with existing goal-reaching methods outperforms a variety of prior methods on open-sourced visual goal-reaching tasks and that Skew-Fit enables a real-world robot to learn to open a door, entirely from scratch, from pixels, and without any manually-designed reward function.

1. Introduction

Reinforcement learning (RL) provides an appealing formalism for automated learning of behavioral skills, but separately learning every potentially useful skill becomes prohibitively time consuming, both in terms of the experience

^{*}Equal contribution ¹University of California, Berkeley. Correspondence to: Vitchyr H. Pong <vitchyr@eecs.berkeley.edu>.




Figure 1. Left: Robot learning to open a door with Skew-Fit, without any task reward. Right: Samples from a goal distribution when using (a) uniform and (b) Skew-Fit sampling. When used as goals, the diverse samples from Skew-Fit encourage the robot to practice opening the door more frequently.

required for the agent and the effort required for the user to design reward functions for each behavior. What if we could instead design an unsupervised RL algorithm that automatically explores the environment and iteratively distills this experience into general-purpose policies that can accomplish new user-specified tasks at test time?

In the absence of any prior knowledge, an effective exploration scheme is one that visits as many states as possible, allowing a policy to autonomously prepare for user-specified tasks that it might see at test time. We can formalize the problem of visiting as many states as possible as one of maximizing the *state entropy* $\mathcal{H}(\mathbf{S})$ under the current policy.² Unfortunately, optimizing this objective alone does not result in a policy that can solve new tasks: it only knows how to maximize state entropy. In other words, to develop principled unsupervised RL algorithms that result in useful policies, maximizing $\mathcal{H}(\mathbf{S})$ is not enough. We need a mechanism that allows us to reuse the resulting policy to achieve new tasks at test-time.

We argue that this can be accomplished by performing *goal-directed exploration*: a policy should autonomously visit as many states as possible, but after autonomous exploration, a user should be able to reuse this policy by giving it a goal \mathbf{G} that corresponds to a state that it must reach. While not all test-time tasks can be expressed as reaching a goal state, a wide range of tasks can be represented in this way. Mathematically, the goal-conditioned policy should minimize the

²We consider the distribution over terminal states in a finite horizon task and believe this work can be extended to infinite horizon stationary distributions.

conditional entropy over the states given a goal, $\mathcal{H}(\mathbf{S} \mid \mathbf{G})$, so that there is little uncertainty over its state given a commanded goal. This objective provides us with a principled way to train a policy to explore all states (maximize $\mathcal{H}(\mathbf{S})$) such that the state that is reached can be determined by commanding goals (minimize $\mathcal{H}(\mathbf{S} \mid \mathbf{G})$).

Directly optimizing this objective is in general intractable, since it requires optimizing the entropy of the marginal state distribution, $\mathcal{H}(\mathbf{S})$. However, we can sidestep this issue by noting that the objective is the mutual information between the state and the goal, $I(\mathbf{S}; \mathbf{G})$, which can be written as:

$$\mathcal{H}(\mathbf{S}) - \mathcal{H}(\mathbf{S} \mid \mathbf{G}) = I(\mathbf{S}; \mathbf{G}) = \mathcal{H}(\mathbf{G}) - \mathcal{H}(\mathbf{G} \mid \mathbf{S}). \quad (1)$$

Equation 1 thus gives an equivalent objective for an unsupervised RL algorithm: the agent should set diverse goals, maximizing $\mathcal{H}(\mathbf{G})$, and learn how to reach them, minimizing $\mathcal{H}(\mathbf{G} \mid \mathbf{S})$.

While learning to reach goals is the typical objective studied in goal-conditioned RL (Kaelbling, 1993; Andrychowicz et al., 2017), setting goals that have maximum diversity is crucial for effectively learning to reach all possible states. Acquiring such a maximum-entropy goal distribution is challenging in environments with complex, high-dimensional state spaces, where even knowing which states are valid presents a major challenge. For example, in image-based domains, a uniform goal distribution requires sampling uniformly from the set of realistic images, which in general is unknown a priori.

Our paper makes the following contributions. First, we propose a principled objective for unsupervised RL, based on Equation 1. While a number of prior works ignore the $\mathcal{H}(\mathbf{G})$ term, we argue that jointly optimizing the entire quantity is needed to develop effective exploration. Second, we present a general algorithm called Skew-Fit and prove that under regularity conditions Skew-Fit learns a sequence of generative models that converges to a uniform distribution over the goal space, even when the set of valid states is unknown (e.g., as in the case of images). Third, we describe a concrete implementation of Skew-Fit and empirically demonstrate that this method achieves state of the art results compared to a large number of prior methods for goal reaching with visually indicated goals, including a real-world manipulation task, which requires a robot to learn to open a door from scratch in about five hours, directly from images, and without any manually-designed reward function.

2. Problem Formulation

To ensure that an unsupervised reinforcement learning agent learns to reach all possible states in a controllable way,

we maximize the mutual information between the state \mathbf{S} and the goal \mathbf{G} , $I(\mathbf{S}; \mathbf{G})$, as stated in Equation 1. This section discusses how to optimize Equation 1 by splitting the optimization into two parts: minimizing $\mathcal{H}(\mathbf{G} \mid \mathbf{S})$ and maximizing $\mathcal{H}(\mathbf{G})$.

2.1. Minimizing $\mathcal{H}(\mathbf{G} \mid \mathbf{S})$: Goal-Conditioned Reinforcement Learning

Standard RL considers a Markov decision process (MDP), which has a state space \mathcal{S} , action space \mathcal{A} , and unknown dynamics $p(\mathbf{s}_{t+1} \mid \mathbf{s}_t, \mathbf{a}_t) : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \mapsto [0, +\infty)$. Goal-conditioned RL also includes a goal space \mathcal{G} . For simplicity, we will assume in our derivation that the goal space matches the state space, such that $\mathcal{G} = \mathcal{S}$, though the approach extends trivially to the case where \mathcal{G} is a hand-specified subset of \mathcal{S} , such as the global XY position of a robot. A goal-conditioned policy $\pi(\mathbf{a} \mid \mathbf{s}, \mathbf{g})$ maps a state $\mathbf{s} \in \mathcal{S}$ and goal $\mathbf{g} \in \mathcal{S}$ to a distribution over actions $\mathbf{a} \in \mathcal{A}$, and its objective is to reach the goal, i.e., to make the current state equal to the goal.

Goal-reaching can be formulated as minimizing $\mathcal{H}(\mathbf{G} \mid \mathbf{S})$, and many practical goal-reaching algorithms (Kaelbling, 1993; Lillicrap et al., 2016; Schaul et al., 2015; Andrychowicz et al., 2017; Nair et al., 2018; Pong et al., 2018; Florensa et al., 2018a) can be viewed as approximations to this objective by observing that the optimal goal-conditioned policy will deterministically reach the goal, resulting in a conditional entropy of zero: $\mathcal{H}(\mathbf{G} \mid \mathbf{S}) = 0$. See Appendix E for more details. Our method may thus be used in conjunction with any of these prior goal-conditioned RL methods in order to jointly minimize $\mathcal{H}(\mathbf{G} \mid \mathbf{S})$ and maximize $\mathcal{H}(\mathbf{G})$.

2.2. Maximizing $\mathcal{H}(\mathbf{G})$: Setting Diverse Goals

We now turn to the problem of setting diverse goals or, mathematically, maximizing the entropy of the goal distribution $\mathcal{H}(\mathbf{G})$. Let $U_{\mathcal{S}}$ be the uniform distribution over \mathcal{S} , where we assume \mathcal{S} has finite volume so that the uniform distribution is well-defined. Let $q_{\phi}^{\mathbf{G}}$ be the goal distribution from which goals \mathbf{G} are sampled, parameterized by ϕ . Our goal is to maximize the entropy of $q_{\phi}^{\mathbf{G}}$, which we write as $\mathcal{H}(\mathbf{G})$. Since the maximum entropy distribution over \mathcal{S} is the uniform distribution $U_{\mathcal{S}}$, maximizing $\mathcal{H}(\mathbf{G})$ may seem as simple as choosing the uniform distribution to be our goal distribution: $q_{\phi}^{\mathbf{G}} = U_{\mathcal{S}}$. However, this requires knowing the uniform distribution over valid states, which may be difficult to obtain when \mathcal{S} is a subset of \mathbb{R}^n , for some n . For example, if the states correspond to images viewed through a robot’s camera, \mathcal{S} corresponds to the (unknown) set of valid images of the robot’s environment, while \mathbb{R}^n corresponds to all possible arrays of pixel values of a particular size. In such environments, sampling from the uniform distribution \mathbb{R}^n is unlikely to correspond to a valid image of the real world.

Sampling uniformly from \mathcal{S} would require knowing the set of all possible valid images, which we assume the agent does not know when starting to explore the environment.

While we cannot sample arbitrary states from \mathcal{S} , we can sample states by performing goal-directed exploration. To derive and analyze our method, we introduce a simple model of this process: a goal $\mathbf{G} \sim q_\phi^G$ is sampled from the goal distribution q_ϕ^G , and then the goal-conditioned policy π attempts to achieve this goal, which results in a distribution of terminal states $\mathbf{S} \in \mathcal{S}$. We abstract this entire process by writing the resulting marginal distribution over \mathbf{S} as $p_\phi^S(\mathbf{S}) \triangleq \int_{\mathcal{G}} q_\phi^G(\mathbf{G})p(\mathbf{S} | \mathbf{G})d\mathbf{G}$, where the subscript ϕ indicates that the marginal p_ϕ^S depends indirectly on q_ϕ^G via the goal-conditioned policy π . We assume that p_ϕ^S has full support, which can be accomplished with an epsilon-greedy goal reaching policy in a communicating MDP. We also assume that the entropy of the resulting state distribution $\mathcal{H}(p_\phi^S)$ is no less than the entropy of the goal distribution $\mathcal{H}(q_\phi^G)$. Without this assumption, a policy could ignore the goal and stay in a single state, no matter how diverse and realistic the goals are.³ This simplified model allows us to analyze the behavior of our goal-setting scheme separately from any specific goal-reaching algorithm. We will however show in Section 6 that we can instantiate this approach into a practical algorithm that jointly learns the goal-reaching policy. In summary, our goal is to acquire a maximum-entropy goal distribution q_ϕ^G over valid states \mathcal{S} , while only having access to state samples from p_ϕ^S .

3. Skew-Fit: Learning a Maximum Entropy Goal Distribution

Our method, Skew-Fit, learns a maximum entropy goal distribution q_ϕ^G using samples collected from a goal-conditioned policy. We analyze the algorithm and show that Skew-Fit maximizes the goal distribution entropy, and present a practical instantiation for unsupervised deep RL.

3.1. Skew-Fit Algorithm

To learn a uniform distribution over *valid* goal states, we present a method that iteratively increases the entropy of a generative model q_ϕ^G . In particular, given a generative model $q_{\phi_t}^G$ at iteration t , we want to train a new generative model, $q_{\phi_{t+1}}^G$ that has higher entropy. While we do not know the set of valid states \mathcal{S} , we could sample states $\mathbf{s}_n \stackrel{\text{iid}}{\sim} p_{\phi_t}^S$ using the goal-conditioned policy, and use the samples to train $q_{\phi_{t+1}}^G$. However, there is no guarantee that this would increase the entropy of $q_{\phi_{t+1}}^G$.

³ Note that this assumption does **not** require that the entropy of p_ϕ^S is strictly larger than the entropy of the goal distribution, q_ϕ^G .

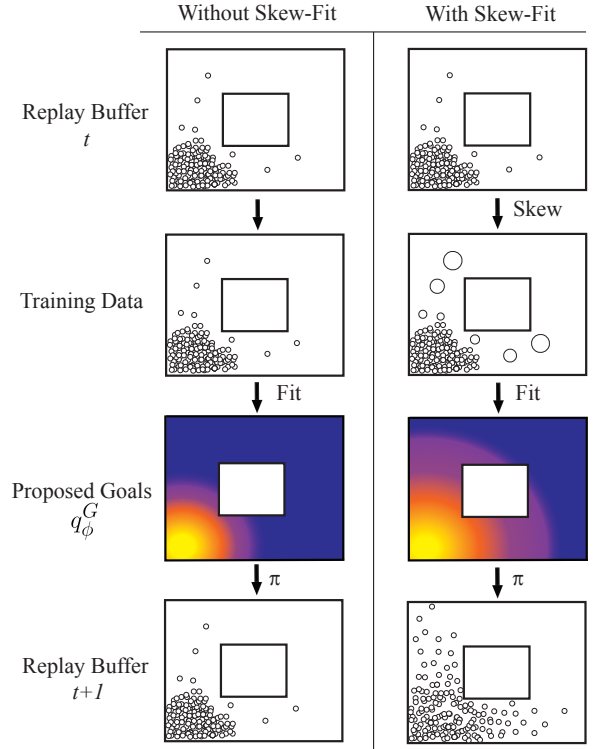


Figure 2. Our method, Skew-Fit, samples goals for goal-conditioned RL. We sample states from our replay buffer, and give more weight to rare states. We then train a generative model $q_{\phi_{t+1}}^G$ with the weighted samples. By sampling new states with goals proposed from this new generative model, we obtain a higher entropy state distribution in the next iteration.

The intuition behind our method is simple: rather than fitting a generative model to these samples \mathbf{s}_n , we *skew* the samples so that rarely visited states are given more weight. See Figure 2 for a visualization of this process. How should we skew the samples if we want to maximize the entropy of $q_{\phi_{t+1}}^G$? If we had access to the density of each state, $p_{\phi_t}^S(\mathbf{S})$, then we could simply weight each state by $1/p_{\phi_t}^S(\mathbf{S})$. We could then perform maximum likelihood estimation (MLE) for the uniform distribution by using the following importance sampling (IS) loss to train ϕ_{t+1} :

$$\begin{aligned} \mathcal{L}(\phi) &= \mathbb{E}_{\mathbf{S} \sim U_{\mathcal{S}}} [\log q_\phi^G(\mathbf{S})] \\ &= \mathbb{E}_{\mathbf{S} \sim p_{\phi_t}^S} \left[\frac{U_{\mathcal{S}}(\mathbf{S})}{p_{\phi_t}^S(\mathbf{S})} \log q_\phi^G(\mathbf{S}) \right] \\ &\propto \mathbb{E}_{\mathbf{S} \sim p_{\phi_t}^S} \left[\frac{1}{p_{\phi_t}^S(\mathbf{S})} \log q_\phi^G(\mathbf{S}) \right] \end{aligned}$$

where we use the fact that the uniform distribution $U_{\mathcal{S}}(\mathbf{S})$ has constant density for all states in \mathcal{S} . However, computing this density $p_{\phi_t}^S(\mathbf{S})$ requires marginalizing out the MDP dynamics, which requires an accurate model of both the dynamics and the goal-conditioned policy.

We avoid needing to model the entire MDP process by approximating $p_{\phi_t}^S(\mathbf{S})$ with our previous learned generative model: $p_{\phi_t}^S(\mathbf{S}) \approx q_{\phi_t}^G(\mathbf{S})$. We therefore weight each state by the following weight function

$$w_{t,\alpha}(\mathbf{S}) \triangleq q_{\phi_t}^G(\mathbf{S})^\alpha, \quad \alpha < 0. \quad (2)$$

where α is a hyperparameter that controls how heavily we weight each state. If our approximation $q_{\phi_t}^G$ is exact, we can choose $\alpha = -1$ and recover the exact IS procedure described above. If $\alpha = 0$, then this skew step has no effect. By choosing intermediate values of α , we trade off the reliability of our estimate $q_{\phi_t}^G(\mathbf{S})$ with the speed at which we want to increase the goal distribution entropy.

Variance Reduction As described, this procedure relies on IS, which can have high variance, particularly if $q_{\phi_t}^G(\mathbf{S}) \approx 0$. We therefore choose a class of generative models where the probabilities are prevented from collapsing to zero, as we will describe in Section 4 where we provide generative model details. To further reduce the variance, we train $q_{\phi_{t+1}}^G$ with sampling importance resampling (SIR) (Rubin, 1988) rather than IS. Rather than sampling from $p_{\phi_t}^S$ and weighting the update from each sample by $w_{t,\alpha}$, SIR explicitly defines a skewed empirical distribution as

$$p_{\text{skewed}_t}(\mathbf{s}) \triangleq \frac{1}{Z_\alpha} w_{t,\alpha}(\mathbf{s}) \delta(\mathbf{s} \in \{\mathbf{s}_n\}_{n=1}^N) \quad (3)$$

$$Z_\alpha = \sum_{n=1}^N w_{t,\alpha}(\mathbf{s}_n), \quad \mathbf{s}_n \stackrel{\text{iid}}{\sim} p_{\phi_t}^S,$$

where δ is the indicator function and Z_α is the normalizing coefficient. We note that computing Z_α adds little computational overhead, since all of the weights already need to be computed. We then fit the generative model at the next iteration $q_{\phi_{t+1}}^G$ to p_{skewed_t} using standard MLE. We found that using SIR resulted in significantly lower variance than IS. See Appendix B.2 for this comparison.

Goal Sampling Alternative Because $q_{\phi_{t+1}}^G \approx p_{\text{skewed}_t}$, at iteration $t + 1$, one can sample goals from either $q_{\phi_{t+1}}^G$ or p_{skewed_t} . Sampling goals from p_{skewed_t} may be preferred if sampling from the learned generative model $q_{\phi_{t+1}}^G$ is computationally or otherwise challenging. In either case, one still needs to train the generative model $q_{\phi_t}^G$ to create p_{skewed_t} . In our experiments, we found that both methods perform well.

Summary Overall, Skew-Fit collects states from the environment and resamples each state in proportion to Equation 2 so that low-density states are resampled more often. Skew-Fit is shown in Figure 2 and summarized in Algorithm 1. We now discuss conditions under which Skew-Fit converges to the uniform distribution.

Algorithm 1 Skew-Fit

- 1: **for** Iteration $t = 1, 2, \dots$ **do**
 - 2: Collect N states $\{\mathbf{s}_n\}_{n=1}^N$ by sampling goals from $q_{\phi_t}^G$ (or $p_{\text{skewed}_{t-1}}$) and running goal-conditioned policy.
 - 3: Construct skewed distribution p_{skewed_t} (Equation 2 and Equation 3).
 - 4: Fit $q_{\phi_{t+1}}^G$ to skewed distribution p_{skewed_t} using MLE.
 - 5: **end for**
-

3.2. Skew-Fit Analysis

This section provides conditions under which $q_{\phi_t}^G$ converges in the limit to the uniform distribution over the state space \mathcal{S} . We consider the case where $N \rightarrow \infty$, which allows us to study the limit behavior of the goal distribution p_{skewed_t} . Our most general result is stated as follows:

Lemma 3.1. *Let \mathcal{S} be a compact set. Define the set of distributions $\mathcal{Q} = \{p : \text{support}(p) \subseteq \mathcal{S}\}$. Let $\mathcal{F} : \mathcal{Q} \mapsto \mathcal{Q}$ be continuous with respect to the pseudometric $d_{\mathcal{H}}(p, q) \triangleq |\mathcal{H}(p) - \mathcal{H}(q)|$ and $\mathcal{H}(\mathcal{F}(p)) \geq \mathcal{H}(p)$ with equality if and only if p is the uniform probability distribution on \mathcal{S} , denoted as $U_{\mathcal{S}}$. Define the sequence of distributions $P = (p_1, p_2, \dots)$ by starting with any $p_1 \in \mathcal{Q}$ and recursively defining $p_{t+1} = \mathcal{F}(p_t)$. The sequence P converges to $U_{\mathcal{S}}$ with respect to $d_{\mathcal{H}}$. In other words, $\lim_{t \rightarrow \infty} |\mathcal{H}(p_t) - \mathcal{H}(U_{\mathcal{S}})| \rightarrow 0$.*

Proof. See Appendix Section A.1. □

We will apply Lemma 3.1 to be the map from p_{skewed_t} to $p_{\text{skewed}_{t+1}}$ to show that p_{skewed_t} converges to $U_{\mathcal{S}}$. If we assume that the goal-conditioned policy and generative model learning procedure are well behaved (i.e., the maps from $q_{\phi_t}^G$ to $p_{\phi_t}^S$ and from p_{skewed_t} to $q_{\phi_{t+1}}^G$ are continuous), then to apply Lemma 3.1, we only need to show that $\mathcal{H}(p_{\text{skewed}_t}) \geq \mathcal{H}(p_{\phi_t}^S)$ with equality if and only if $p_{\phi_t}^S = U_{\mathcal{S}}$. For the simple case when $q_{\phi_t}^G = p_{\phi_t}^S$ identically at each iteration, we prove the convergence of Skew-Fit true for any value of $\alpha \in [-1, 0)$ in Appendix A.3. However, in practice, $q_{\phi_t}^G$ only approximates $p_{\phi_t}^S$. To address this more realistic situation, we prove the following result:

Lemma 3.2. *Given two distribution $p_{\phi_t}^S$ and $q_{\phi_t}^G$ where $p_{\phi_t}^S \ll q_{\phi_t}^G$ ⁴ and*

$$\text{Cov}_{\mathbf{S} \sim p_{\phi_t}^S} [\log p_{\phi_t}^S(\mathbf{S}), \log q_{\phi_t}^G(\mathbf{S})] > 0, \quad (4)$$

define the p_{skewed_t} as in Equation 3 and take $N \rightarrow \infty$. Let $\mathcal{H}_\alpha(\alpha)$ be the entropy of p_{skewed_t} for a fixed α . Then there exists a constant $a < 0$ such that for all $\alpha \in [a, 0)$,

$$\mathcal{H}(p_{\text{skewed}_t}) = \mathcal{H}_\alpha(\alpha) > \mathcal{H}(p_{\phi_t}^S).$$

⁴ $p \ll q$ means that p is absolutely continuous with respect to q , i.e. $p(\mathbf{s}) = 0 \implies q(\mathbf{s}) = 0$.

Proof. See Appendix Section A.2. \square

This lemma tells us that our generative model $q_{\phi_t}^G$ does not need to exactly fit the sampled states. Rather, we merely need the log densities of $q_{\phi_t}^G$ and $p_{\phi_t}^S$ to be correlated, which we expect to happen frequently with an accurate goal-conditioned policy, since $p_{\phi_t}^S$ is the set of states seen when trying to reach goals from $q_{\phi_t}^G$. In this case, if we choose negative values of α that are small enough, then the entropy of p_{skewed_t} will be higher than that of $p_{\phi_t}^S$. Empirically, we found that α values as low as $\alpha = -1$ performed well.

In summary, p_{skewed_t} converges to U_S under certain assumptions. Since we train each generative model $q_{\phi_{t+1}}^G$ by fitting it to p_{skewed_t} with MLE, $q_{\phi_t}^G$ will also converge to U_S .

4. Training Goal-Conditioned Policies with Skew-Fit

Thus far, we have presented Skew-Fit assuming that we have access to a goal-reaching policy, allowing us to separately analyze how we can maximize $\mathcal{H}(\mathbf{G})$. However, in practice we do not have access to such a policy, and this section discusses how we concurrently train a goal-reaching policy.

Maximizing $I(\mathbf{S}; \mathbf{G})$ can be done by simultaneously performing Skew-Fit and training a goal-conditioned policy to minimize $\mathcal{H}(\mathbf{G} | \mathbf{S})$, or, equivalently, maximize $-\mathcal{H}(\mathbf{G} | \mathbf{S})$. Maximizing $-\mathcal{H}(\mathbf{G} | \mathbf{S})$ requires computing the density $\log p(\mathbf{G} | \mathbf{S})$, which may be difficult to compute without strong modeling assumptions. However, for any distribution q , the following lower bound on $-\mathcal{H}(\mathbf{G} | \mathbf{S})$:

$$\begin{aligned} -\mathcal{H}(\mathbf{G} | \mathbf{S}) &= \mathbb{E}_{(\mathbf{G}, \mathbf{S}) \sim q} [\log q(\mathbf{G} | \mathbf{S})] + D_{\text{KL}}(p \| q) \\ &\geq \mathbb{E}_{(\mathbf{G}, \mathbf{S}) \sim q} [\log q(\mathbf{G} | \mathbf{S})], \end{aligned}$$

where D_{KL} denotes Kullback–Leibler divergence as discussed by Barber & Agakov (2004). Thus, to minimize $\mathcal{H}(\mathbf{G} | \mathbf{S})$, we train a policy to maximize the reward

$$r(\mathbf{S}, \mathbf{G}) = \log q(\mathbf{G} | \mathbf{S}).$$

The RL algorithm we use is reinforcement learning with imagined goals (RIG) (Nair et al., 2018), though in principle any goal-conditioned method could be used. RIG is an efficient off-policy goal-conditioned method that solves vision-based RL problems in a learned latent space. In particular, RIG fits a β -VAE (Higgins et al., 2017) and uses it to encode observations and goals into a latent space, which it uses as the state representation. RIG also uses the β -VAE to compute rewards, $\log q(\mathbf{G} | \mathbf{S})$. Unlike RIG, we use the goal distribution from Skew-Fit to sample goals for exploration and for relabeling goals during training (Andrychowicz et al., 2017). Since RIG already trains a generative model over states, we reuse this β -VAE for the generative

model q_{ϕ}^G of Skew-Fit. To make the most use of the data, we train q_{ϕ}^G on all visited state rather than only the terminal states, which we found to work well in practice. To prevent the estimated state likelihoods from collapsing to zero, we model the posterior of the β -VAE as a multivariate Gaussian distribution with a fixed variance and only learn the mean. We summarize RIG and provide details for how we combine Skew-Fit and RIG in Appendix C.4 and describe how we estimate the likelihoods given the β -VAE in Appendix C.1.

5. Related Work

Many prior methods in the goal-conditioned reinforcement learning literature focus on training goal-conditioned policies and assume that a goal distribution is available to sample from during exploration (Kaelbling, 1993; Schaul et al., 2015; Andrychowicz et al., 2017; Pong et al., 2018), or use a heuristic to design a non-parametric (Colas et al., 2018b; Warde-Farley et al., 2018; Florensa et al., 2018a) or parametric (P  r   et al., 2018; Nair et al., 2018) goal distribution based on previously visited states. These methods are largely complementary to our work: rather than proposing a better method for training goal-reaching policies, we propose a principled method for maximizing the entropy of a goal sampling distribution, $\mathcal{H}(\mathbf{G})$, such that these policies cover a wide range of states.

Our method learns without any task rewards, directly acquiring a policy that can be reused to reach user-specified goals. This stands in contrast to exploration methods that modify the reward based on state visitation frequency (Bellemare et al., 2016; Ostrovski et al., 2017; Tang et al., 2017; Chen-tanez et al., 2005; Lopes et al., 2012; Stadie et al., 2016; Pathak et al., 2017; Burda et al., 2018; 2019; Mohamed & Rezende, 2015; Tang et al., 2017; Fu et al., 2017). While these methods can also be used without a task reward, they provide no mechanism for distilling the knowledge gained from visiting diverse states into flexible policies that can be applied to accomplish new goals at test-time: their policies visit novel states, and they quickly forget about them as other states become more novel. Similarly, methods that provably maximize state entropy without using goal-directed exploration (Hazan et al., 2019) or methods that define new rewards to capture measures of intrinsic motivation (Mohamed & Rezende, 2015) and reachability (Savinov et al., 2018) do not produce reusable policies.

Other prior methods extract reusable skills in the form of latent-variable-conditioned policies, where latent variables are interpreted as options (Sutton et al., 1999) or abstract skills (Hausman et al., 2018; Gupta et al., 2018b; Eysenbach et al., 2019; Gupta et al., 2018a; Florensa et al., 2017). The resulting skills are diverse, but have no grounded interpretation, while Skew-Fit policies can be used immediately after unsupervised training to reach diverse user-specified goals.




Figure 3. Illustrative example of Skew-Fit on a 2D navigation task. (Left) Visited state plot for Skew-Fit with $\alpha = -1$ and uniform sampling, which corresponds to $\alpha = 0$. (Right) The entropy of the goal distribution per iteration, mean and standard deviation for 9 seeds. Entropy is calculated via discretization onto an 11×11 grid. Skew-Fit steadily increases the state entropy, reaching full coverage over the state space.

Some prior methods propose to choose goals based on heuristics such as learning progress (Baranes & Oudeyer, 2012; Veeriah et al., 2018; Colas et al., 2018a), how off-policy the goal is (Nachum et al., 2018), level of difficulty (Florensa et al., 2018b), or likelihood ranking (Zhao & Tresp, 2019). In contrast, our approach provides a principled framework for optimizing a concrete and well-motivated exploration objective, can provably maximize this objective under regularity assumptions, and empirically outperforms many of these prior work (see Section 6).

6. Experiments

Our experiments study the following questions: (1) Does Skew-Fit empirically result in a goal distribution with increasing entropy? (2) Does Skew-Fit improve exploration for goal-conditioned RL? (3) How does Skew-Fit compare to prior work on choosing goals for vision-based, goal-conditioned RL? (4) Can Skew-Fit be applied to a real-world, vision-based robot task?

Does Skew-Fit Maximize Entropy? To see the effects of Skew-Fit on goal distribution entropy in isolation of learning a goal-reaching policy, we study an idealized example where the policy is a near-perfect goal-reaching policy. The environment consists of four rooms (Sutton et al., 1999). At the beginning of an episode, the agent begins in the bottom-right room and samples a goal from the goal distribution $q_{\phi_t}^G$. To simulate stochasticity of the policy and environment, we add a Gaussian noise with standard deviation of 0.06 units to this goal, where the entire environment is 11×11 units. The policy reaches the state that is closest to this noisy goal and inside the rooms, giving us a state sample s_n for training $q_{\phi_t}^G$. Due to the relatively small noise, the agent cannot rely on this stochasticity to explore the different rooms and must instead learn to set goals that are progressively farther and farther from the initial state. We compare multiple values of α , where $\alpha = 0$ corresponds to not using Skew-Fit. The β -VAE hyperparameters used to train $q_{\phi_t}^G$ are given in Appendix C.2. As seen in Figure 3, sampling uniformly from previous experience ($\alpha = 0$) to set goals results in a policy




Figure 4. (Left) Ant navigation environment. (Right) Evaluation on reaching target XY position. We show the mean and standard deviation of 6 seeds. Skew-Fit significantly outperforms prior methods on this exploration task.

that primarily sets goal near the initial state distribution. In contrast, Skew-Fit results in quickly learning a high entropy, near-uniform distribution over the state space.

Exploration with Skew-Fit We next evaluate Skew-Fit while concurrently learning a goal-conditioned policy on a task with state inputs, which enables us study exploration performance independently of the challenges with image observations. We evaluate on a task that requires training a simulated quadruped “ant” robot to navigate to different XY positions in a labyrinth, as shown in Figure 4. The reward is the negative distance to the goal XY-position, and additional environment details are provided in Appendix D. This task presents a challenge for goal-directed exploration: the set of valid goals is unknown due to the walls, and random actions do not result in exploring locations far from the start. Thus, Skew-Fit must set goals that meaningfully explore the space while simultaneously learning to reach those goals.

We use this domain to compare Skew-Fit to a number of existing goal-sampling methods. We compare to the relabeling scheme described in the hindsight experience replay (labeled **HER**). We compare to curiosity-driven prioritization (**Ranked-Based Priority**) (Zhao et al., 2019), a variant of HER that samples goals for relabeling based on their ranked likelihoods. Florensa et al. (2018b) samples goals from a GAN based on the difficulty of reaching the goal. We compare against this method by replacing q_{ϕ}^G with the GAN and label it **AutoGoal GAN**. We also compare to the non-parametric goal proposal mechanism proposed by (Warde-Farley et al., 2018), which we label **DISCERN-g**. Lastly, to demonstrate the difficulty of the exploration challenge in these domains, we compare to **#-Exploration** (Tang et al., 2017), an exploration method that assigns bonus rewards based on the novelty of new states. We train the goal-conditioned policy for each method using soft actor critic (SAC) (Haarnoja et al., 2018). Implementation details of SAC and the prior works are given in Appendix C.3.

We see in Figure 4 that Skew-Fit is the only method that makes significant progress on this challenging labyrinth locomotion task. The prior methods on goal-sampling primarily set goals close to the start location, while the extrinsic exploration reward in #-Exploration dominated the goal-reaching reward. These results demonstrate that Skew-Fit




Figure 5. We evaluate on these continuous control tasks, from left to right: *Visual Door*, a door opening task; *Visual Pickup*, a picking task; *Visual Pusher*, a pushing task; and *Real World Visual Door*, a real world door opening task. All tasks are solved from images and without any task-specific reward. See Appendix D for details.

accelerates exploration by setting diverse goals in tasks with unknown goal spaces.

Vision-Based Continuous Control Tasks We now evaluate Skew-Fit on a variety of image-based continuous control tasks, where the policy must control a robot arm using only image observations, there is no state-based or task-specific reward, and Skew-Fit must directly set image goals. We test our method on three different image-based simulated continuous control tasks released by the authors of RIG (Nair et al., 2018): *Visual Door*, *Visual Pusher*, and *Visual Pickup*. These environments contain a robot that can open a door, push a puck, and lift up a ball to different configurations, respectively. To our knowledge, these are the only goal-conditioned, vision-based continuous control environments that are publicly available and experimentally evaluated in prior work, making them a good point of comparison. See Figure 5 for visuals and Appendix C for environment details. The policies are trained in a completely unsupervised manner, without access to any prior information about the image-space or any pre-defined goal-sampling distribution. To evaluate their performance, we sample goal images from a uniform distribution over valid states and report the agent’s final distance to the corresponding simulator states (e.g., distance of the object to the target object location), but the agent never has access to this true uniform distribution nor the ground-truth state information during training. While this evaluation method is only practical in simulation, it provides us with a quantitative measure of a policy’s ability to reach a broad coverage of goals in a vision-based setting.

We compare Skew-Fit to a number of existing methods on this domain. First, we compare to the methods described in the previous experiment (HER, Rank-Based Priority, #-Exploration, AutoGoal GAN, and DISCERN-g). These methods that we compare to were developed in non-vision, state-based environments. To ensure a fair comparison

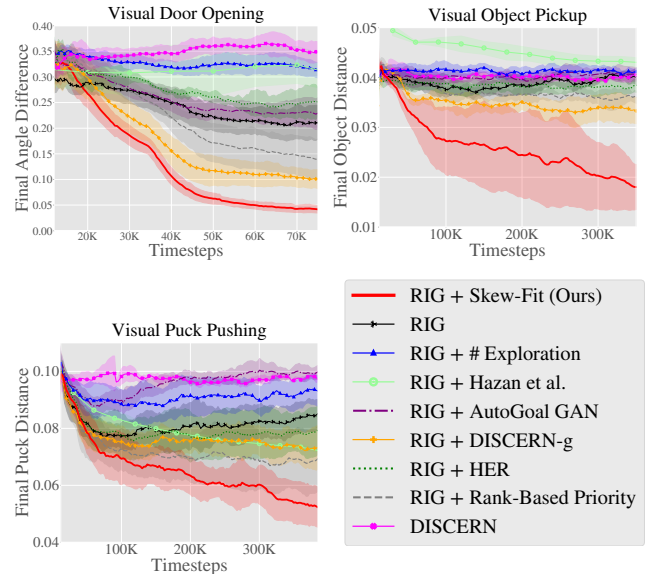


Figure 6. Learning curves for simulated continuous control tasks. Lower is better. We show the mean and standard deviation of 6 seeds and smooth temporally across 50 epochs within each seed. Skew-Fit consistently outperforms RIG and various prior methods. See text for description of each method.

across methods, we combine these prior methods with a policy trained using RIG. We additionally compare to Hazan et al. (2019), an exploration method that assigns bonus rewards based on the likelihood of a state (labeled Hazan et al.). Next, we compare to RIG without Skew-Fit. Lastly, we compare to DISCERN (Warde-Farley et al., 2018), a vision-based method which uses a non-parametric clustering approach to sample goals and an image discriminator to compute rewards.

We see in Figure 6 that Skew-Fit significantly outperforms prior methods both in terms of task performance and sample complexity. The most common failure mode for prior methods is that the goal distributions collapse, resulting in the agent learning to reach only a fraction of the state space, as shown in Figure 1. For comparison, additional samples of q_ϕ^G when trained with and without Skew-Fit are shown in Appendix B.3. Those images show that without Skew-Fit, q_ϕ^G produces a small, non-diverse distribution for each environment: the object is in the same place for pickup, the puck is often in the starting position for pushing, and the door is always closed. In contrast, Skew-Fit proposes goals where the object is in the air and on the ground, where the puck positions are varied, and the door angle changes.

We can see the effect of these goal choices by visualizing more example rollouts for RIG and Skew-Fit. These visuals, shown in Figure 14 in Appendix B.3, show that RIG only learns to reach states close to the initial position, while Skew-Fit learns to reach the entire state space. For a quantitative comparison, Figure 7 shows the cumulative total exploration




Figure 7. Cumulative total pickups during exploration for each method. Prior methods fail to pay attention to the object: the rate of pickups hardly increases past the first 100 thousand timesteps. In contrast, after seeing the object picked up a few times, Skew-Fit practices picking up the object more often by sampling the appropriate exploration goals.

pickups for each method. From the graph, we see that many methods have a near-constant rate of object lifts throughout all of training. Skew-Fit is the only method that significantly increases the rate at which the policy picks up the object during exploration, suggesting that only Skew-Fit sets goals that encourage the policy to interact with the object.

Real-World Vision-Based Robotic Manipulation We also demonstrate that Skew-Fit scales well to the real world with a door opening task, *Real World Visual Door*, as shown in Figure 5. While a number of prior works have studied RL-based learning of door opening (Kalakrishnan et al., 2011; Chebotar et al., 2017), we demonstrate the first method for autonomous learning of door opening without a user-provided, task-specific reward function. As in simulation, we do not provide any goals to the agent and simply let it interact with the door, without any human guidance or reward signal. We train two agents using RIG and RIG with Skew-Fit. Every seven and a half minutes of interaction time, we evaluate on 5 goals and plot the cumulative successes for each method. Unlike in simulation, we cannot easily measure the difference between the policy’s achieved and desired door angle. Instead, we visually denote a binary success/failure for each goal based on whether the last state in the trajectory achieves the target angle. As Figure 8 shows, standard RIG only starts to open the door after five hours of training. In contrast, Skew-Fit learns to occasionally open the door after three hours of training and achieves a near-perfect success rate after five and a half hours of interaction. Figure 8 also shows examples of successful trajectories from the Skew-Fit policy, where we see that the policy can reach a variety of user-specified goals. These results demonstrate that Skew-Fit is a promising technique for solving real world tasks without any human-provided reward function. Videos of Skew-Fit solving this task and the simulated tasks can be viewed on our website.⁵

⁵<https://sites.google.com/view/skew-fit>

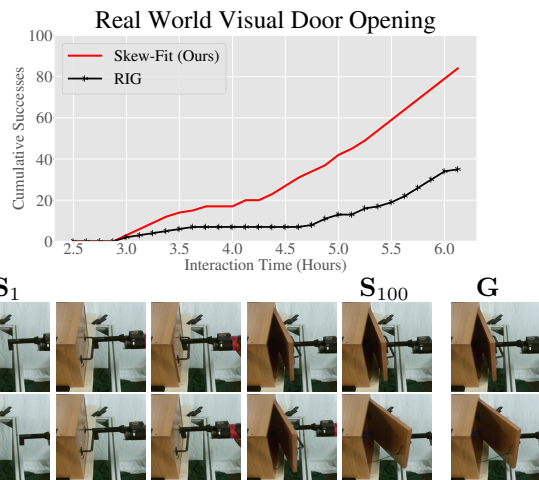


Figure 8. (Top) Learning curve for Real World Visual Door. Skew-Fit results in considerable sample efficiency gains over RIG on this real-world task. (Bottom) Each row shows the Skew-Fit policy starting from state S_1 and reaching state S_{100} while pursuing goal G . Despite being trained from only images without any user-provided goals during training, the Skew-Fit policy achieves the goal image provided at test-time, successfully opening the door.

Additional Experiments To study the sensitivity of Skew-Fit to the hyperparameter α , we sweep α across the values $[-1, -0.75, -0.5, -0.25, 0]$ on the simulated image-based tasks. The results are in Appendix B and demonstrate that Skew-Fit works across a large range of values for α , and $\alpha = -1$ consistently outperform $\alpha = 0$ (i.e. outperforms no Skew-Fit). Additionally, Appendix C provides a complete description our method hyperparameters, including network architecture and RL algorithm hyperparameters.

7. Conclusion

We presented a formal objective for self-supervised goal-directed exploration, allowing researchers to quantify and compare progress when designing algorithms that enable agents to autonomously learn. We also presented Skew-Fit, an algorithm for training a generative model to approximate a uniform distribution over an initially unknown set of valid states, using data obtained via goal-conditioned reinforcement learning, and our theoretical analysis gives conditions under which Skew-Fit converges to the uniform distribution. When such a model is used to choose goals for exploration and to relabeling goals for training, the resulting method results in much better coverage of the state space, enabling our method to explore effectively. Our experiments show that when we concurrently train a goal-reaching policy using self-generated goals, Skew-Fit produces quantifiable improvements on simulated robotic manipulation tasks, and can be used to learn a door opening skill to reach a 95% success rate directly on a real-world robot, without any human-provided reward supervision.

8. Acknowledgement

This research was supported by Berkeley DeepDrive, Huawei, ARL DCIST CRA W911NF-17-2-0181, NSF IIS-1651843, and the Office of Naval Research, as well as Amazon, Google, and NVIDIA. We thank Aviral Kumar, Carlos Florensa, Aurick Zhou, Nilesh Tripuraneni, Vickie Ye, Dibya Ghosh, Coline Devin, Rowan McAllister, John D. Co-Reyes, various members of the Berkeley Robotic AI & Learning (RAIL) lab, and anonymous reviewers for their insightful discussions and feedback.

References

- Andrychowicz, M., Wolski, F., Ray, A., Schneider, J., Fong, R., Welinder, P., McGrew, B., Tobin, J., Abbeel, P., and Zaremba, W. Hindsight Experience Replay. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- Baranes, A. and Oudeyer, P.-Y. Active Learning of Inverse Models with Intrinsically Motivated Goal Exploration in Robots. *Robotics and Autonomous Systems*, 61(1):49–73, 2012. doi: 10.1016/j.robot.2012.05.008.
- Barber, D. and Agakov, F. V. Information maximization in noisy channels: A variational approach. In *Advances in Neural Information Processing Systems*, pp. 201–208, 2004.
- Bellemare, M., Srinivasan, S., Ostrovski, G., Schaul, T., Saxton, D., and Munos, R. Unifying count-based exploration and intrinsic motivation. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 1471–1479, 2016.
- Burda, Y., Edwards, H., Storkey, A., and Klimov, O. Exploration by random network distillation. *arXiv preprint arXiv:1810.12894*, 2018.
- Burda, Y., Edwards, H., Pathak, D., Storkey, A., Darrell, T., and Efros, A. A. Large-scale study of curiosity-driven learning. In *International Conference on Learning Representations (ICLR)*, 2019.
- Chebotaev, Y., Kalakrishnan, M., Yahya, A., Li, A., Schaal, S., and Levine, S. Path integral guided policy search. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3381–3388. IEEE, 2017.
- Chentanez, N., Barto, A. G., and Singh, S. P. Intrinsically motivated reinforcement learning. In *Advances in neural information processing systems*, pp. 1281–1288, 2005.
- Colas, C., Fournier, P., Sigaud, O., and Oudeyer, P. CURIOS: intrinsically motivated multi-task, multi-goal reinforcement learning. *CoRR*, abs/1810.06284, 2018a.
- Colas, C., Sigaud, O., and Oudeyer, P.-Y. Gep-pg: Decoupling exploration and exploitation in deep reinforcement learning algorithms. *International Conference on Machine Learning (ICML)*, 2018b.
- Eysenbach, B., Gupta, A., Ibarz, J., and Levine, S. Diversity is All You Need: Learning Skills without a Reward Function. In *International Conference on Learning Representations (ICLR)*, 2019.
- Florensa, C., Duan, Y., and Abbeel, P. Stochastic neural networks for hierarchical reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2017.
- Florensa, C., Degraeve, J., Heess, N., Springenberg, J. T., and Riedmiller, M. Self-supervised Learning of Image Embedding for Continuous Control. In *Workshop on Inference to Control at NeurIPS*, 2018a.
- Florensa, C., Held, D., Geng, X., and Abbeel, P. Automatic Goal Generation for Reinforcement Learning Agents. In *International Conference on Machine Learning (ICML)*, 2018b.
- Fu, J., Co-Reyes, J. D., and Levine, S. EX 2 : Exploration with Exemplar Models for Deep Reinforcement Learning. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- Fujimoto, S., van Hoof, H., and Meger, D. Addressing Function Approximation Error in Actor-Critic Methods. In *International Conference on Machine Learning (ICML)*, 2018.
- Gupta, A., Eysenbach, B., Finn, C., and Levine, S. Unsupervised meta-learning for reinforcement learning. *CoRR*, abs:1806.04640, 2018a.
- Gupta, A., Mendonca, R., Liu, Y., Abbeel, P., and Levine, S. Meta-Reinforcement Learning of Structured Exploration Strategies. In *Advances in Neural Information Processing Systems (NIPS)*, 2018b.
- Haarnoja, T., Zhou, A., Hartikainen, K., Tucker, G., Ha, S., Tan, J., Kumar, V., Zhu, H., Gupta, A., Abbeel, P., and Levine, S. Soft actor-critic algorithms and applications. *CoRR*, abs/1812.05905, 2018.
- Hausman, K., Springenberg, J. T., Wang, Z., Heess, N., and Riedmiller, M. Learning an Embedding Space for Transferable Robot Skills. In *International Conference on Learning Representations (ICLR)*, pp. 1–16, 2018.
- Hazan, E., Kakade, S. M., Singh, K., and Soest, A. V. Provably efficient maximum entropy exploration. *International Conference on Machine Learning (ICML)*, 2019.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. β -vae: Learning basic visual concepts with a constrained variational framework. *International Conference on Learning Representations (ICLR)*, 2017.
- Kaelbling, L. P. Learning to achieve goals. In *International Joint Conference on Artificial Intelligence (IJCAI)*, volume vol.2, pp. 1094 – 8, 1993.
- Kalakrishnan, M., Righetti, L., Pastor, P., and Schaal, S. Learning force control policies for compliant manipulation. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 4639–4644. IEEE, 2011.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2016. ISBN 0-7803-3213-X. doi: 10.1613/jair.301.
- Lopes, M., Lang, T., Toussaint, M., and Oudeyer, P.-Y. Exploration in model-based reinforcement learning by empirically estimating learning progress. In *Advances in Neural Information Processing Systems*, pp. 206–214, 2012.

- Mohamed, S. and Rezende, D. J. Variational information maximisation for intrinsically motivated reinforcement learning. In *Advances in neural information processing systems*, pp. 2125–2133, 2015.
- Nachum, O., Brain, G., Gu, S., Lee, H., and Levine, S. Data-Efficient Hierarchical Reinforcement Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- Nair, A., Pong, V., Dalal, M., Bahl, S., Lin, S., and Levine, S. Visual Reinforcement Learning with Imagined Goals. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- Nielsen, F. and Nock, R. Entropies and cross-entropies of exponential families. In *Image Processing (ICIP), 2010 17th IEEE International Conference on*, pp. 3621–3624. IEEE, 2010.
- Ostrovski, G., Bellemare, M. G., Oord, A., and Munos, R. Count-based exploration with neural density models. In *International Conference on Machine Learning (ICML)*, pp. 2721–2730, 2017.
- Pathak, D., Agrawal, P., Efros, A. A., and Darrell, T. Curiosity-Driven Exploration by Self-Supervised Prediction. In *International Conference on Machine Learning (ICML)*, pp. 488–489. IEEE, 2017.
- Péré, A., Forestier, S., Sigaud, O., and Oudeyer, P.-Y. Unsupervised Learning of Goal Spaces for Intrinsically Motivated Goal Exploration. In *International Conference on Learning Representations (ICLR)*, 2018.
- Pong, V., Gu, S., Dalal, M., and Levine, S. Temporal Difference Models: Model-Free Deep RL For Model-Based Control. In *International Conference on Learning Representations (ICLR)*, 2018.
- Rubin, D. B. Using the sir algorithm to simulate posterior distributions. *Bayesian statistics*, 3:395–402, 1988.
- Savinov, N., Raichuk, A., Marinier, R., Vincent, D., Pollefeys, M., Lillicrap, T., and Gelly, S. Episodic curiosity through reachability. *arXiv preprint arXiv:1810.02274*, 2018.
- Schaul, T., Horgan, D., Gregor, K., and Silver, D. Universal Value Function Approximators. In *International Conference on Machine Learning (ICML)*, pp. 1312–1320, 2015. ISBN 9781510810587.
- Stadie, B. C., Levine, S., and Abbeel, P. Incentivizing Exploration In Reinforcement Learning With Deep Predictive Models. In *International Conference on Learning Representations (ICLR)*, 2016.
- Sutton, R. S., Precup, D., and Singh, S. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2):181–211, 1999.
- Tang, H., Houthoofd, R., Foote, D., Stooke, A., Chen, X., Duan, Y., Schulman, J., De Turck, F., and Abbeel, P. #Exploration: A Study of Count-Based Exploration for Deep Reinforcement Learning. In *Neural Information Processing Systems (NIPS)*, 2017.
- Todorov, E., Erez, T., and Tassa, Y. MuJoCo: A physics engine for model-based control. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5026–5033, 2012. ISBN 9781467317375. doi: 10.1109/IROS.2012.6386109.
- Veeriah, V., Oh, J., and Singh, S. Many-goals reinforcement learning. *arXiv preprint arXiv:1806.09605*, 2018.
- Warde-Farley, D., de Wiele, T. V., Kulkarni, T., Ionescu, C., Hansen, S., and Mnih, V. Unsupervised control through non-parametric discriminative rewards. *CoRR*, abs/1811.11359, 2018.
- Zhao, R. and Tresp, V. Curiosity-driven experience prioritization via density estimation. *CoRR*, abs/1902.08039, 2019.
- Zhao, R., Sun, X., and Tresp, V. Maximum entropy-regularized multi-goal reinforcement learning. In *International Conference on Machine Learning*, pp. 7553–7562, 2019.