

Goal-Conditioned Reinforcement Learning with Imagined Subgoals

Elliot Chane-Sane¹ Cordelia Schmid¹ Ivan Laptev¹

Abstract

Goal-conditioned reinforcement learning endows an agent with a large variety of skills, but it often struggles to solve tasks that require more temporally extended reasoning. In this work, we propose to incorporate imagined subgoals into policy learning to facilitate learning of complex tasks. Imagined subgoals are predicted by a separate high-level policy, which is trained simultaneously with the policy and its critic. This high-level policy predicts intermediate states halfway to the goal using the value function as a reachability metric. We don't require the policy to reach these subgoals explicitly. Instead, we use them to define a prior policy, and incorporate this prior into a KL-constrained policy iteration scheme to speed up and regularize learning. Imagined subgoals are used during policy learning, but not during test time, where we only apply the learned policy. We evaluate our approach on complex robotic navigation and manipulation tasks and show that it outperforms existing methods by a large margin.

1. Introduction

An intelligent agent aims at solving tasks of varying horizons in an environment. It must be able to identify how the different tasks intertwine with each other and leverage behaviors learned by solving simpler tasks to efficiently master more complex tasks. For instance, once a legged robot has learned how to walk in every direction in a simulated maze environment, it could use these behaviors to efficiently learn to navigate in this environment. Goal-conditioned reinforcement learning (RL) (Kaelbling, 1993; Schaul et al., 2015; Andrychowicz et al., 2017) defines each task by the desired goal and, in principle, could learn a wide range of skills. In practice, however, reinforcement learning struggles to perform temporally extended reasoning.

¹Inria, École normale supérieure, CNRS, PSL Research University, 75005 Paris, France. Correspondence to: Elliot Chane-Sane <elliot.chane-sane@inria.fr>.




Figure 1. Illustration of the KL-regularized policy learning using imagined subgoals. (Left): The policy fails to reach a distant goal, yet it can reach a closer subgoal. Our approach automatically generates imagined subgoals for a task and uses such subgoals to direct the policy search during training. (Right): At test time, the resulting flat policy can reach arbitrarily distant goals *without* relying on subgoals.

Hierarchical methods have proven effective for learning temporal abstraction in long-horizon problems (Dayan & Hinton, 1993; Wiering & Schmidhuber, 1997; Sutton et al., 1999; Dietterich, 2000). In the goal-conditioned setting, a high-level controller typically finds an appropriate sequence of subgoals that can be more easily followed by a low-level policy (Nachum et al., 2018; Gupta et al., 2019). When chosen appropriately, these subgoals effectively decompose a complex task into easier tasks. However, hierarchical methods are often unstable to train (Nachum et al., 2018) and rely on appropriate temporal design choices.

In this work, we propose to use subgoals to improve the goal-conditioned policy. Instead of reaching the subgoals explicitly, our method builds on the following intuition. If the current policy can readily reach a subgoal, it could provide guidance for reaching more distant goals, as illustrated in Figure 1. We apply this idea to all possible state and goal pairs of the environment. This self-supervised approach progressively extends the horizon of the agent throughout training. At the end of training, the resulting flat policy does not require access to subgoals and can reach distant goals in its environment.

Our method does not require a full sequence of subgoals, but predicts subgoals that are halfway to the goal, such

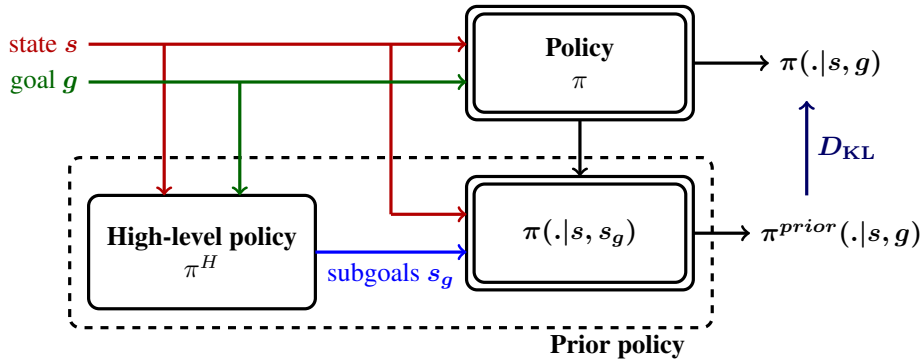


Figure 2. Overview over our Reinforcement learning with Imagined Subgoals (RIS) approach. During policy training, the policy π is constrained to stay close to the prior policy π^{prior} through KL-regularization. We define the prior policy π^{prior} as the distribution of actions required to reach intermediate subgoals s_g of the task. Given the initial state s and the goal state g , the subgoals are generated by the high-level policy π^H . Note that the high-level policy and subgoals are only used during training of the target policy π . At test time we directly use π to generate appropriate actions.

that reaching them corresponds to a lower level of temporal abstraction than reaching the goal. To handle this subgoal prediction problem, we simultaneously train a separate high-level policy operating in the state space and use the value function of the goal-conditioned policy as a relative measure of distance between states (Eysenbach et al., 2019; Nasiriany et al., 2019).

To incorporate subgoals into policy learning, we introduce a prior policy defined as the distribution over actions required to reach intermediate subgoals (Figure 2). When using appropriate subgoals that are easier to reach, this prior policy corresponds to an initial guess to reach the goal. Accordingly, we leverage a policy iteration with an additional Kullback-Leibler (KL) constraint scheme towards this prior policy. This adequately encourages the policy to adapt the simpler behaviors associated with reaching these subgoals to master more complex goal-reaching tasks in the environment. Because these subgoals are *not* actively pursued when interacting with the actual environment but only used to accelerate policy learning, we refer to them as *imagined subgoals*. Imagined subgoals and intermediate subgoals have equal meaning in this paper.

Our method, Reinforcement learning with Imagined Subgoals (RIS), builds upon off-policy actor-critic approaches for continuous control and additionally learns a high-level policy that updates its predictions according to the current goal-reaching capabilities of the policy. Our approach is general, solves temporally extended goal-reaching tasks with sparse rewards, and self-supervises its training by choosing appropriate imagined subgoals to accelerate its learning.

In summary, our contributions are threefold: (1) we propose a method for predicting subgoals which decomposes a goal-conditioned task into easier subtasks; (2) we incorporate these subgoals into policy learning through a KL-regularized policy iteration scheme with a specific choice of prior policy;

and (3) we show that our approach greatly accelerates policy learning on a set of simulated robotics environments that involve motor control and temporally extended reasoning.¹

2. Related Work

Goal-conditioned reinforcement learning has been addressed by a number of methods (Kaelbling, 1993; Schaul et al., 2015; Andrychowicz et al., 2017; Veeriah et al., 2018; Pong et al., 2018; Nair et al., 2018; Zhao et al., 2019; Pitis et al., 2020; Eysenbach et al., 2020). Given the current state and the goal, the resulting policies predict action sequences that lead towards the desired goal. Hindsight experience replay (HER) (Kaelbling, 1993; Andrychowicz et al., 2017) is often used to improve the robustness and sample efficiency of goal-reaching policies. While in theory such policies can address any goal-reaching task, they often fail to solve temporally extended problems in practice (Levy et al., 2019; Nachum et al., 2018).

Long-horizon tasks can be addressed by hierarchical reinforcement learning (Dayan & Hinton, 1993; Wiering & Schmidhuber, 1997; Dietterich, 2000; Levy et al., 2019; Vezhnevets et al., 2017). Such methods often design high-level policies that operate at a coarser time scale and control execution of low-level policies. To address goal-conditioned settings, high-level policies can be learned to iteratively predict a sequence of intermediate subgoals. Such subgoals can then be used as targets for low-level policies (Nachum et al., 2018; Gupta et al., 2019; Nair & Finn, 2020). As an alternative to the iterative planning, other methods generate sequences of subgoals with a divide-and-conquer approach (Jurgenson et al., 2020; Parascandolo et al., 2020; Pertsch et al., 2020b). While hierarchical RL methods can

¹Code is available on the project webpage <https://www.di.ens.fr/willow/research/ris/>.

better address long-horizon tasks, the joint learning of high-level and low-level policies may imply instabilities (Nachum et al., 2018). Similar to previous hierarchical RL methods, we use subgoals to decompose long-horizon tasks into simpler problems. Our subgoals, however, are only used during policy learning to guide and accelerate the search of the non-hierarchical policy.

Several recent RL methods use the value function of goal-reaching policies to measure distances between states and to plan sequences of subgoals (Nasiriany et al., 2019; Eysenbach et al., 2019; Zhang et al., 2020). In particular, LEAP (Nasiriany et al., 2019) uses the value function of TDM policies (Pong et al., 2018) and optimizes sequences of appropriate subgoals at test time. Similar to previous methods, we use the value function as a distance measure between states. Our method, however, avoids expensive test-time optimization of subgoals. We also experimentally compare our method with LEAP and demonstrate improvements. Moreover, we show that our approach can benefit from recent advances in learning representations for reinforcement learning from pixels on a vision-based robotic manipulation task (Kostrikov et al., 2020; Laskin et al., 2020).

Many approaches cast reinforcement learning as a probabilistic inference problem where the optimal policy should match a probability distribution in a graphical model defined by the reward and the environment dynamics (Toussaint, 2009; Kappen et al., 2012; Levine, 2018). Several methods optimize an objective incorporating the divergence between the target policy and a prior policy. The prior policy can be fixed (Haarnoja et al., 2017; Abdolmaleki et al., 2018b; Haarnoja et al., 2018a; Pertsch et al., 2020a) or learned jointly with the policy (Teh et al., 2017; Galashov et al., 2019; Tirumala et al., 2019). While previous work imposes explicit priors e.g., in the multi-task and transfer learning settings (Teh et al., 2017; Galashov et al., 2019; Tirumala et al., 2019), we constrain our policy search by the prior distribution implicitly defined by subgoals produced by the high-level policy.

Behavior priors have often been used to avoid value overestimation for out-of-distribution actions in offline reinforcement learning (Fujimoto et al., 2018; Wu et al., 2019; Kumar et al., 2019; Siegel et al., 2020; Nair et al., 2020; Wang et al., 2020). Similar to these methods, we constrain our high-level policy to avoid predicting subgoals outside of the valid state distribution.

Finally, our work shares similarities with guided policy search methods (Levine & Koltun, 2013; Levine & Abbeel, 2014; Levine et al., 2016) which alternate between generating expert trajectories using trajectory optimization and improving the learned policy. In contrast, our policy search is guided by subgoals produced by a high-level policy.

3. Method

Our method builds on the following key observation. If an action a is well-suited for approaching an intermediate subgoal s_g from state s , it should also be a good choice for approaching the final goal g of the same task. We assume that reaching subgoals s_g is simpler than reaching more distant goals g . Hence, we adopt a self-supervised strategy and use the subgoal-reaching policy $\pi(\cdot|s, s_g)$ as a guidance when learning the goal-reaching policy $\pi(\cdot|s, g)$. We denote our approach as Reinforcement learning with Imagined Subgoals (RIS) and present its overview in Figures 1 and 2.

To implement the idea of RIS, we first introduce a high-level policy π^H predicting imagined subgoals s_g , as described in Section 3.2. Section 3.3 presents the regularized learning of the target policy π using subgoals. The joint training of π and π^H is summarized in Section 3.4. Before describing our technical contributions, we present the actor-critic paradigm used by RIS in Section 3.1 below.

3.1. Goal-Conditioned Actor-Critic

We consider a discounted, infinite-horizon, goal-conditioned Markov decision process, with states $s \in \mathcal{S}$, goals $g \in \mathcal{G}$, actions $a \in \mathcal{A}$, reward function $r(s, a, g)$, dynamics $p(s'|s, a)$ and discount factor γ . The objective of a goal-conditioned RL agent is to maximize the expected discounted return

$$J(\pi) = \mathbb{E}_{g \sim \rho_g, \tau \sim d^\pi(\cdot|g)} \left[\sum_t \gamma^t r(s_t, a_t, g) \right]$$

under the distribution

$$d^\pi(\tau|g) = \rho_0(s_0) \prod_t \pi(a_t|s_t, g) p(s_{t+1}|s_t, a_t)$$

induced by the policy π and the initial state and goal distribution. The policy $\pi(\cdot|s, g)$ in this work generates a distribution over continuous actions a conditioned on state s and goal g . Many algorithms rely on the appropriate learning of the goal-conditioned action-value function Q^π and the value function V^π defined as $Q^\pi(s, a, g) = \mathbb{E}_{s_0=s, a_0=a, \tau \sim d^\pi(\cdot|g)} [\sum_t \gamma^t r(s_t, a_t, g)]$ and $V^\pi(s, g) = \mathbb{E}_{a \sim \pi(\cdot|s, g)} Q^\pi(s, a, g)$.

In this work we assume states and goals to co-exist in the same space, i.e. $\mathcal{S} = \mathcal{G}$, where each state can be considered as a potential goal. Moreover, we set the reward r to -1 for all actions until the policy reaches the goal.

We follow the standard off-policy actor-critic paradigm (Silver et al., 2014; Heess et al., 2015; Mnih et al., 2016; Fujimoto et al., 2018; Haarnoja et al., 2018a). Experience, consisting of single transition tuples (s_t, a_t, s_{t+1}, g) , is collected by the policy in a replay buffer D . Actor-critic algorithms maximize return by alternating between policy evaluation and policy improvement. During the policy evaluation phase, a critic $Q^\pi(s, a, g)$ estimates the action-value




Figure 3. Given initial states s and goal states g , our subgoals are states on the middle of the path from s to g . We measure distances between states by the value function $|V^\pi(s_1, s_2)|$ corresponding to the current policy π . We obtain a distribution of subgoals from the high-level policy $s_g \sim \pi^H(\cdot|s, g)$. We use subgoals *only* at the training to regularize and accelerate the policy search.

function of the current policy π by minimizing the Bellman error with respect to the Q-function parameters ϕ_k :

$$Q_{\phi_{k+1}} = \arg \min_{\phi} \frac{1}{2} \mathbb{E}_{(s_t, a_t, s_{t+1}, g) \sim D} [y_t - Q_{\phi}(s_t, a_t, g)]^2 \quad (1)$$

with the target value

$$y_t = r(s_t, a_t, g) + \gamma \mathbb{E}_{a_{t+1} \sim \pi(\cdot|s_t, g)} Q_{\phi_k}(s_{t+1}, a_{t+1}, g).$$

During the policy improvement phase, the actor π is typically updated such that the expected value of the current Q-function Q^π , or alternatively the advantage $A^\pi(s, a, g) = Q^\pi(s, a, g) - V^\pi(s, g)$, under π is maximized:

$$\pi_{\theta_{k+1}} = \arg \max_{\theta} \mathbb{E}_{(s, g) \sim D, a \sim \pi_{\theta}(\cdot|s, g)} [Q^\pi(s, a, g)]. \quad (2)$$

Reaching distant goals using delayed rewards may require expensive policy search. To accelerate the training, we propose to direct the policy search towards intermediate subgoals of a task. Our method learns to predict appropriate subgoals by the high-level policy π^H . The high-level policy is trained together with the policy π as explained below.

3.2. High-Level Policy

We would like to learn a high-level policy $\pi^H(\cdot|s, g)$ that predicts an appropriate distribution of imagined subgoals s_g conditioned on valid states s and goals g . Our high-level policy is defined in terms of the policy π and relies on the goal-reaching capabilities of π as described next.

Subgoal search with a value function. We note that our choice of the reward function $r = -1$ implies that the norm of the value function $V^\pi(s, g)$ corresponds to an estimate of the expected discounted number of steps required for the policy to reach the goal g from the current state s . We therefore propose to use $|V^\pi(s^i, s^j)|$ as a measure of the distance between any valid states s^i and s^j . Note that this measure depends on the policy π and evolves with the improvements of π during training.

Reaching imagined subgoals of a task should be easier than reaching the final goal. To leverage this assumption for policy learning, we need to find appropriate subgoals. In this work we propose to define subgoals s_g as midpoints on the path from the current state s to the goal g , see Figure 3. More formally, we wish s_g (i) to have equal distance from s and g , and (ii) to minimize the length of the paths from s to s_g and from s_g to g . We can find subgoals that satisfy these constraints by using our distance measure $|V^\pi(s^i, s^j)|$ and minimizing the following cost $C_\pi(s_g|s, g)$

$$C_\pi(s_g|s, g) = \max(|V^\pi(s, s_g)|, |V^\pi(s_g, g)|). \quad (3)$$

However, naively minimizing the cost C_π under the high-level policy distribution, i.e.

$$\pi_{k+1}^H = \arg \min_{\pi^H} \mathbb{E}_{(s, g) \sim D, s_g \sim \pi^H(\cdot|s, g)} [C_\pi(s_g|s, g)], \quad (4)$$

may lead to undesired solutions where the high-level policy samples subgoals outside the valid state distribution $p_s(\cdot)$. Such predictions may, for example, correspond to unfeasible robot poses, unrealistic images, or other adversarial states which have low distance from s and g according to $|V^\pi|$ but are unreachable in practice.

Predicting valid subgoals. To avoid non-valid subgoals, we can additionally encourage the high-level policy to stay close to the valid state distribution $p_s(\cdot)$ using the following KL-regularized objective:

$$\begin{aligned} \pi_{k+1}^H &= \arg \max_{\pi^H} \mathbb{E}_{(s, g) \sim D, s_g \sim \pi^H(\cdot|s, g)} [A^{\pi_k^H}(s_g|s, g)] \\ \text{s.t. } &D_{\text{KL}}(\pi^H(\cdot|s, g) || p_s(\cdot)) \leq \epsilon, \end{aligned} \quad (5)$$

where the advantage

$$A^{\pi_k^H}(s_g|s, g) = \mathbb{E}_{s_g \sim \pi_k^H(\cdot|s, g)} [C_\pi(s_g|s, g)] - C_\pi(s_g|s, g)$$

quantifies the quality of a subgoal s_g against the high-level policy distribution. The Kullback-Leibler divergence term in (5) requires an estimate of the density $p_s(\cdot)$. While estimating the unknown $p_s(\cdot)$ might be challenging, we can obtain samples from this distribution, for example, by randomly sampling states from the replay buffer $s_g \sim D$.

We propose instead to *implicitly* enforce the KL constraint (5). We first note that the analytic solution to (5) can be obtained by enforcing the Karush–Kuhn–Tucker conditions, for which the Lagrangian is:

$$\mathcal{L}(\pi^H, \lambda) = \mathbb{E}_{s_g \sim \pi(\cdot|s, g)} \left[A^{\pi^H}(s_g|s, g) \right] + \lambda (\epsilon - D_{\text{KL}}(\pi(\cdot|s, g) \| p_s(\cdot))).$$

The closed form solution to this problem is then given by:

$$\pi_{\star}^H(s_g|s, g) = \frac{1}{Z(s, g)} p_s(s_g) \exp\left(\frac{1}{\lambda} A^{\pi^H}(s_g|s, g)\right) \quad (6)$$

with the normalizing partition function $Z(s, g) = \int p_s(s_g) \exp\left(\frac{1}{\lambda} A^{\pi^H}(s_g|s, g)\right) ds_g$ (Peters et al., 2010; Rawlik et al., 2012; Abdolmaleki et al., 2018b;a; Nair et al., 2020). We project this solution into the policy space by minimizing the forward KL divergence between our parametric high-level policy π_{ψ}^H and the optimal non parametric solution π_{\star}^H :

$$\begin{aligned} \pi_{\psi_{k+1}}^H &= \arg \min_{\psi} \mathbb{E}_{(s, g) \sim D} D_{\text{KL}}(\pi_{\star}^H(\cdot|s, g) \| \pi_{\psi}^H(\cdot|s, g)) \\ &= \arg \max_{\psi} \mathbb{E}_{(s, g) \sim D, s_g \sim D} \left[\log \pi_{\psi}^H(s_g|s, g) \frac{1}{Z(s, g)} \exp\left(\frac{1}{\lambda} A^{\pi^H}(s_g|s, g)\right) \right], \end{aligned} \quad (7)$$

where λ is a hyperparameter. Conveniently, this policy improvement step corresponds to a weighted maximum likelihood with subgoal candidates obtained from the replay buffer randomly sampled among states visited by the agent in previous episodes. The samples are re-weighted by their corresponding advantage, implicitly constraining the high-level policy to stay close to the valid distribution of states. A more detailed derivation is given in Appendix A.

3.3. Policy Improvement with Imagined Subgoals

Our method builds on the following key insight. If we assume s_g to be an intermediate subgoal on the optimal path from s to g , then the optimal action for reaching g from s should be similar to the optimal action for reaching s_g from s (see Figure 1). We can formalize this using a KL constraint on the policy distribution, conditioned on goals g and s_g :

$$D_{\text{KL}}(\pi(\cdot|s, g) \| \pi(\cdot|s, s_g)) \leq \epsilon.$$

We introduce the non-parametric prior policy $\pi^{\text{prior}}(\cdot|s, g)$ as the distribution over actions that would be chosen by the policy π for reaching subgoals $s_g \sim \pi^H(\cdot|s, g)$ provided by the high-level policy. Given a state s and a goal g , we bootstrap the behavior of the policy at subgoals $s_g \sim \pi^H(\cdot|s, g)$ (see Figure 2):

$$\pi_k^{\text{prior}}(a|s, g) := \mathbb{E}_{s_g \sim \pi^H(\cdot|s, g)} [\pi_{\theta'_k}(a|s, s_g)]. \quad (8)$$

As we assume the subgoals to be easier to reach compared to reaching the final goals, this prior policy provides a good initial guess to constrain the policy search to the most promising actions.

We then propose to leverage a policy iteration scheme with additional KL constraint to shape the policy behavior accordingly. During the policy improvement step, in addition to maximizing the Q-function as in (2), we encourage the policy to stay close to the prior policy through the KL-regularization:

$$\begin{aligned} \pi_{\theta_{k+1}} &= \arg \max_{\theta} \mathbb{E}_{(s, g) \sim D} \mathbb{E}_{a \sim \pi_{\theta}(\cdot|s, g)} \\ &\quad \left[Q^{\pi}(s, a, g) - \alpha D_{\text{KL}}\left(\pi_{\theta}(\cdot|s, g) \| \pi_k^{\text{prior}}(\cdot|s, g)\right) \right], \end{aligned} \quad (9)$$

where α is a hyperparameter.

In practice, we found that using an exponential moving average of the online policy parameters to construct the prior policy is necessary to ensure convergence:

$$\theta'_{k+1} = \tau \theta_k + (1 - \tau) \theta'_k, \quad \tau \in]0, 1[. \quad (10)$$

This ensures that the prior policy produces a more stable target for regularizing the online policy.

3.4. Algorithm Summary

Algorithm 1 RL with imagined subgoals

```

Initialize replay buffer  $D$ 
Initialize  $Q_{\phi}, \pi_{\theta}, \pi_{\psi}^H$ 
for  $k = 1, 2, \dots$  do
    Collect experience in  $D$  using  $\pi_{\theta}$  in the environment
    Sample batch  $(s_t, a_t, r_t, s_{t+1}, g) \sim D$  with HER
    Sample batch of subgoal candidates  $s_g \sim D$ 
    Update  $Q_{\phi}$  using Eq. 1 (Policy Evaluation)
    Update  $\pi_{\psi}^H$  Using Eq. 7 (High-Level Policy Improvement)
    Update  $\pi_{\theta}$  using Eq. 9 (Policy Improvement with Imagined Subgoals)
end for
    
```

We approximate the policy π_{θ} , the Q-function Q_{ϕ} and the high-level policy π_{ψ}^H with neural networks parametrized by θ , ϕ and ψ respectively, and train them jointly using stochastic gradient descent. The Q-function is trained to minimize the Bellman error (1), where we use an exponential moving average of the online Q-function parameters to compute the target value. The high-level policy is a probabilistic neural network whose output parametrizes a Laplace distribution with diagonal variance $\pi_{\psi}^H(\cdot|s, g) = \text{Laplace}(\mu_{\psi}^H(s, g), \Sigma_{\psi}^H(s, g))$ trained to minimize (7). The policy is also a probabilistic network parametrizing a squashed Gaussian distribution with diagonal variance $\pi_{\psi}^H(\cdot|s, g) = \tanh \mathcal{N}(\mu_{\theta}(s, g), \Sigma_{\theta}(s, g))$

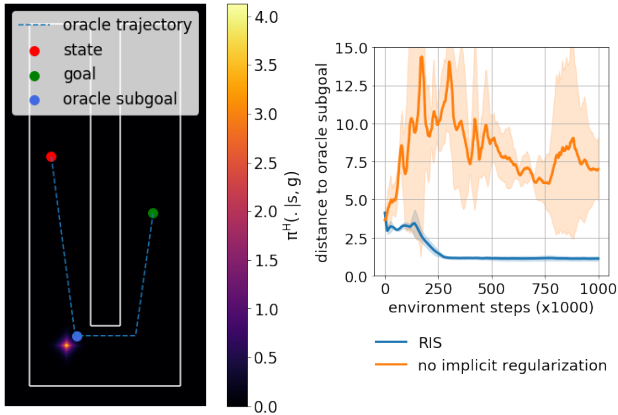


Figure 4. (Left) Heatmap of the subgoal distribution obtained with our high-level policy and the oracle subgoal for a given state and goal for the ant U-maze environment. (Right) Distance between oracle subgoals and subgoals predicted by the high-level policy for RIS and RIS without implicit regularization. The dimensions of the space are 7.5×18 units and the ant has a radius of roughly 0.75 units.

trained to minimize (9). Finally we can approximate π_{prior} using a Monte-Carlo estimate of (8).

The policy π and the high-level policy π^H are trained jointly. As the policy learns to reach more and more distant goals, its value function becomes a better estimate for the distance between states. This allows the high-level policy to propose more appropriate subgoals for a larger set of goals. In turn, as the high-level policy improves, imagined subgoals offer a more relevant supervision to shape the behavior of the policy. This virtuous cycle progressively extends the policy horizon further and further away, allowing more complex tasks to be solved by a single policy.

The full algorithm is summarized in Algorithm 1. We use hindsight experience replay (Andrychowicz et al., 2017) (HER) to improve learning from sparse rewards. Additional implementation details are given in Appendix B.

4. Experiments

In this section we first introduce our experimental setup in Section 4.1. Next, we ablate various design choices of our approach in Section 4.2. We, then, compare RIS to prior work in goal-conditioned reinforcement learning in Section 4.3.

4.1. Experimental Setup

Ant navigation. We evaluate RIS on a set of ant navigation tasks of increasing difficulty, each of which requires temporally extended reasoning. In these environments, the agent observes the joint angles, joint velocity, and center of mass of a quadruped ant robot navigating in a maze. We




Figure 5. Ablation of our method on the Ant U-Maze environment: simple priors that do not incorporate subgoals (*uniform prior*; *moving average prior*); ignoring the effect of out-of-distribution subgoal predictions (*no implicit regularization*); and using oracle subgoals (*with oracle subgoals*). Left: success rate on all configurations throughout training. Right: success rate on the test configurations.

consider four different mazes: a U-shaped maze, a S-shaped maze, a Π -shaped maze and a ω -shaped maze illustrated in Figure 6. The obstacles are unknown to the agent.

During training, initial states and goals are uniformly sampled and the agents are trained to reach any goal in the environment. We evaluate agents in the most extended temporal settings representing the most difficult configurations offered by the environment (see Figure 6). We assess the success rate achieved by these agents, where we define success as the ant being sufficiently close to the goal position measured by x-y Euclidean distance.

Vision-based robotic manipulation. We follow the experimental setup in (Nasiriany et al., 2019) and also consider a vision-based robotic manipulation task where an agent controls a 2 DoF robotic arm from image input and must manipulate a puck positioned on the table (Figure 8a). We define success as the arm and the puck being sufficiently close to their respective desired positions. During training, the initial arm and puck positions and their respective desired positions are uniformly sampled whereas, at test time, we evaluate agents on temporally extended configurations.

These tasks are challenging because they require temporally extended reasoning on top of complex motor control. Indeed, a greedy path towards the goal cannot solve these tasks. We train the agents for 1 million environment steps and average the results over 4 random seeds. Additional details about the environments are given in Appendix C.

Alternative methods. We compare our approach, RIS, to off-policy reinforcement learning methods for goal-reaching tasks. We consider Soft Actor-Critic (SAC) (Haarnoja et al., 2018a) with HER, which trains a policy from scratch by maximizing the entropy regularized objective using the same sparse reward as ours. We also compare to Temporal Difference Models (TDM) (Pong et al., 2018) which trains

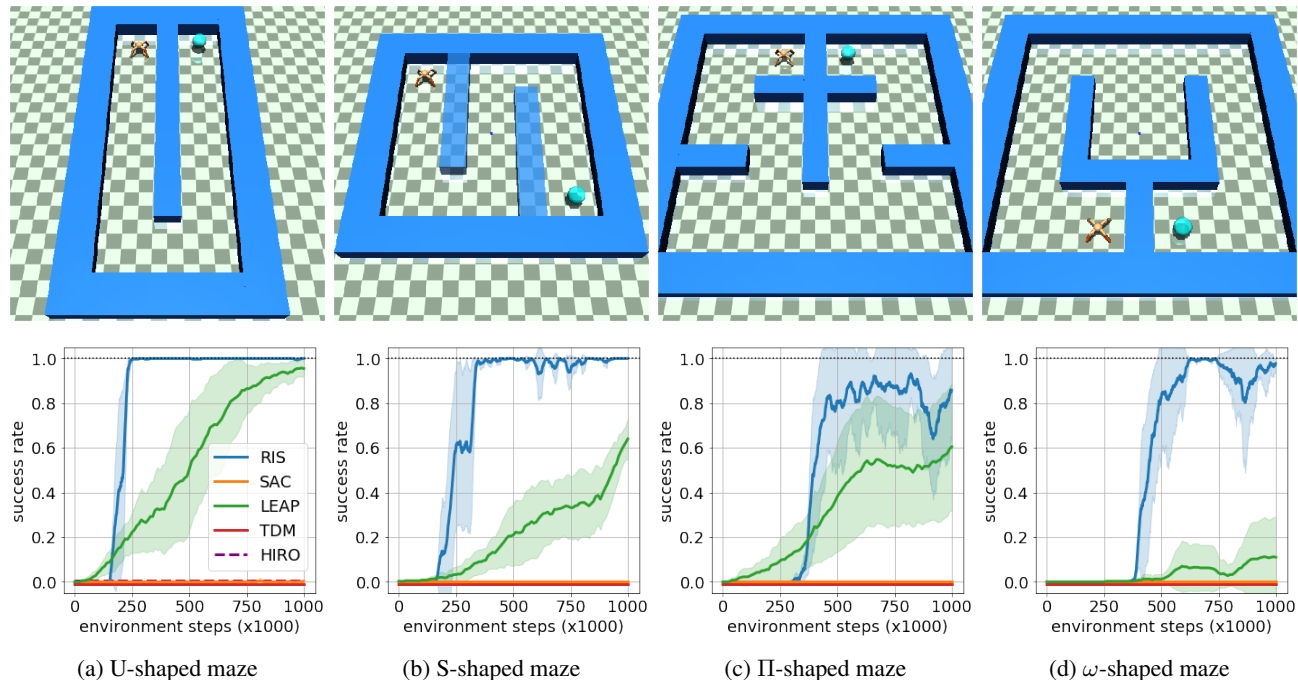


Figure 6. Comparison of RIS to several state-of-the-art methods (bottom row) on 4 different ant navigation tasks. We evaluate the success rate of the agent on the challenging configurations illustrates in the top row, where the ant is located at the initial state and the desired goal location is represented by a cyan sphere.

horizon-aware policies operating under dense rewards in the form of distance to the goal. We chose to evaluate TDMs with a long policy horizon of 600 steps due to the complexity of considered tasks. Furthermore, we compare to Latent Embeddings for Abstracted Planning (LEAP) (Nasiriany et al., 2019), a competitive approach for these environments, which uses a sequence of subgoals that a TDM policy must reach one after the other during inference.

We re-implemented SAC, TDM and LEAP and validated our implementations on the U-shaped ant maze and vision-based robotic manipulation environments.

On the U-shaped Ant maze environment, we additionally report results of HIRO (Nachum et al., 2018), a hierarchical reinforcement learning method with off-policy correction, after 1 million environment steps. The results were copied from Figure 12 in (Nasiriany et al., 2019).

We provide additional details about the hyperparameters used in our experiments in Appendix D.

4.2. Ablative Analysis

We use the Ant U-maze navigation task to conduct ablation studies to validate the role of our design choices.

Evaluating imagined subgoals. To evaluate the quality of the subgoals predicted by our high-level policy, we intro-

duce an oracle subgoal sampling procedure. We plan an oracle trajectory, which corresponds to a point-mass agent navigating in this maze and does not necessarily correspond to the optimal trajectory of an ant. Oracle subgoals correspond to the midpoint of this trajectory between state and goal in x-y location (Figure 4 left).

Figure 4 (left) shows the subgoal distribution predicted by a fully-trained high-level policy for a state and goal located at opposite sides of the U-shaped maze. We can observe that its probability mass is close to the oracle subgoal.

To quantitatively evaluate the quality of imagined subgoals, we measure the x-y Euclidean distance between oracle subgoals and subgoals sampled from the high-level policy throughout training for a set of fixed state and goal tuples randomly sampled in the environment. Figure 4 (right) shows that RIS successfully learns to find subgoals that are coherent with the oracle trajectory during training, despite not having prior knowledge about the environment.

We also assess the importance of using the implicit regularization scheme presented in Section 3.2 which discourages high-level policy predictions to lie outside of the distribution of valid states. We compare our approach against naively optimizing subgoals without regularization (i.e. directly optimize (4)). Figure 4 (right) shows that without implicit regularization, the predicted subgoals significantly diverge




Figure 7. Comparison of different methods for increasingly more difficult tasks on the S-shaped and ω -shaped ant maze. The goals with increasing complexity are numbered starting with 1, where 1 is the closest goal.

from oracle subgoals in x-y location during training. As result, imagined subgoals are not able to properly guide policy learning to solve the task, see Figure 5.

Prior policy with imagined subgoals. We evaluate the importance of incorporating imagined subgoals into the prior policy by comparing to a number of variants. To disentangle the effects of imagined subgoals from the actor-critic architecture used by RIS, we first replace our prior policy with simpler choices of prior distributions that do not incorporate any subgoals: (i) a uniform prior over the actions $\pi_k^{prior} = \mathcal{U}(\mathcal{A})$, which is equivalent to SAC without entropy regularization during policy evaluation, and (ii) a parametric prior policy that uses an exponential moving average of the online policy weights $\pi_k^{prior} = \pi_{\theta_k}$. Figure 5 shows that, while they can learn goal-reaching behaviors on many configurations encountered during training (left), neither of these variants are able to solve the Ant U-maze environment in its most difficult setting (right). We also observe that the agent with a moving average action prior fails to learn (left). This highlights the benefits of incorporating subgoals into policy learning.

Finally, we propose to incorporate the oracle subgoals into our prior policy. We replace the subgoal distributions predicted by our high-level policy with Laplace distributions centered around oracle subgoals. Results in Figure 5 show that RIS with oracle subgoals learns to solve the U-shaped ant maze environment faster than using a high-level policy simultaneously trained by the agent. This experiment highlights the efficiency of our approach to guide policy learning with appropriate subgoals: if we have access to proper subgoals right from the beginning of the training, our approach could leverage them to learn even faster. However, such subgoals are generally not readily available without prior knowledge about the environment. Thus, we introduce a high-level policy training procedure which determines appropriate subgoals without any supervision.

4.3. Comparison to the State of the Art

Ant navigation. Figure 6 compares RIS to the alternative methods introduced in section 4.1 for the four ant navigation environments. For all considered mazes RIS significantly outperforms prior methods in terms of sample efficiency, often requiring less than 500 thousand environment interactions to solve the mazes in their most challenging initial state and goal configurations.

LEAP makes progress on these navigation tasks, but requires significantly more environment interactions. This comparison shows the effectiveness of our approach, which uses subgoals to guide the policy rather than reaching them sequentially as done by LEAP. While SAC manages to learn goal-reaching behaviors, as we will see later in Figure 7, it fails to solve the environments in their most challenging configurations. The comparison to SAC highlights the benefits of using our informed prior policy compared to methods assuming a uniform action prior. On the U-shaped maze environment, HIRO similarly fails to solve the task within one million environment interactions. Furthermore, we observe that TDMs fails to learn due to the sparsity of the reward.

Figure 7 evaluates how SAC, LEAP and RIS perform for varying task horizons in the S-shaped and ω -shaped mazes. Starting from an initial state located at the edge of the mazes, we sample goals at locations which require an increasing number of environment steps to be reached. Figure 7 reports results for RIS, LEAP and SAC after having been trained for 1 million steps. While the performances of LEAP and SAC degrades as the planning horizon increases, RIS consistently solves configurations of increasing complexity.

These results demonstrate that RIS manages to solve complex navigation tasks despite relying on a flat policy at inference. In contrast, LEAP performs less well, despite a significantly more expensive planning of subgoals during inference.

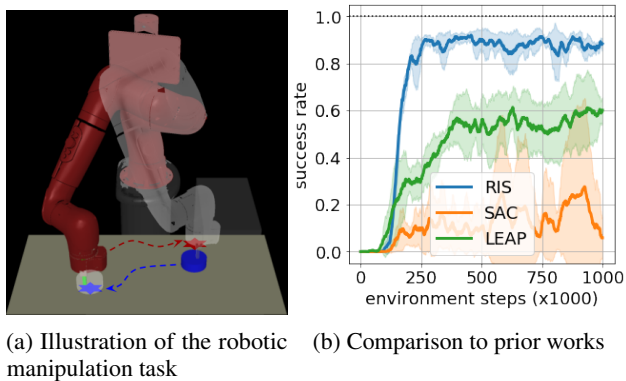


Figure 8. Robotic manipulation environment: (a) illustration of the task; (b) results of our method compared to LEAP and SAC.

Vision-based robotic manipulation. While the ant navigation experiments demonstrate the performances of RIS on environments with low-dimension state spaces, we also show how our method can be applied to vision-based robotic manipulation tasks. Our approach takes images of the current and desired configurations as input. Input images are passed through an image encoder, a convolutional neural network shared between the policy, high-level policy and Q-function. The encoder is only updated when training the Q-function during policy evaluation and is fixed otherwise during policy improvement and high-level policy improvement. Instead of generating subgoals in the high-dimensional image space, the high-level policy therefore operates in the learned compact image representation of the encoder. Following recent works on reinforcement learning from images, we augment image observations with random translations (Kostrikov et al., 2020; Laskin et al., 2020). We found that using such data augmentation was important for training image-based RIS and SAC policies. Moreover, we found that using a lower learning rate for the policy was necessary to stabilize training. Additional implementation details on the image encoding are given in Appendix B.

We compare our approach against LEAP and SAC in Figure 8b. RIS achieves a higher success rate than LEAP whereas SAC fails most of the time to solve the manipulation task consistently enough in the temporally extended configuration used for evaluation on the vision-based robotic manipulation task. Moreover, RIS and SAC only requires a single forward pass through their image encoder and actor network at each time step when interacting with the environment, whereas LEAP depends in addition upon an expensive planning of image subgoals.

Figure 9 visualizes the imagined subgoals of the high-level policy. Once the RIS agent is fully trained, we separately train a decoder to reconstruct image observations from their learned representations. Given observations of the current




Figure 9. Image reconstruction of an imagined subgoal (middle) given the current state (left) and the desired goal (right) on a temporally extended configuration used for evaluation (top) and a random configuration (bottom).

state and the desired goal, we then predict the representation of an imagined subgoal with the high-level policy and generate the corresponding image using the decoder. Figure 9 shows that subgoals predicted by the high-level policy are natural intermediate states halfway to the desired goal on this manipulation task. For example, for the test configuration (Figure 9 top), the high-level policy prediction corresponds to a configuration where the arm has reached the right side of the puck and is pushing it towards its desired position.

Additional reconstructions of imagined subgoals for different initial state and goal configurations in the robotic manipulation environment are given in Appendix E.

5. Conclusion

We introduced RIS, a goal-conditioned reinforcement learning method that imagines possible subgoals in a self-supervised fashion and uses them to facilitate training. We propose to use the value function of the goal-reaching policy to train a high-level policy operating in the state space. We then use imagined subgoals to define a prior policy and incorporate this prior into policy learning. Experimental results on challenging simulated navigation and vision-based manipulation environments show that our proposed method greatly accelerates learning of temporally extended tasks and outperforms competing approaches.

While our approach makes use of subgoals to facilitate policy search, future work could explore how to use them to obtain better Q-value estimates. Future work could also improve exploration by using imagined subgoals to encourage the policy to visit all potential states.

References

- Abdolmaleki, A., Springenberg, J. T., Degraeve, J., Bohez, S., Tassa, Y., Belov, D., Heess, N., and Riedmiller, M. Relative entropy regularized policy iteration. *arXiv preprint arXiv:1812.02256*, 2018a.
- Abdolmaleki, A., Springenberg, J. T., Tassa, Y., Munos, R., Heess, N., and Riedmiller, M. A. Maximum a posteriori policy optimisation. In *International Conference on Learning Representations*, 2018b.
- Andrychowicz, M., Crow, D., 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*, 2017.
- Dayan, P. and Hinton, G. E. Feudal reinforcement learning. In *Advances in Neural Information Processing Systems*, 1993.
- Dietterich, T. G. Hierarchical reinforcement learning with the maxq value function decomposition. *Journal of artificial intelligence research*, 13, 2000.
- Eysenbach, B., Salakhutdinov, R., and Levine, S. Search on the replay buffer: Bridging planning and reinforcement learning. In *Advances in Neural Information Processing Systems*, 2019.
- Eysenbach, B., Salakhutdinov, R., and Levine, S. C-learning: Learning to achieve goals via recursive classification. *arXiv preprint arXiv:2011.08909*, 2020.
- Fujimoto, S., van Hoof, H., and Meger, D. Addressing function approximation error in actor-critic methods. In *International Conference on Machine Learning*, 2018.
- Galashov, A., Jayakumar, S. M., Hasenclever, L., Tirumala, D., Schwarz, J., Desjardins, G., Czarnecki, W. M., Teh, Y. W., Pascanu, R., and Heess, N. Information asymmetry in kl-regularized RL. In *International Conference on Learning Representations*, 2019.
- Gupta, A., Kumar, V., Lynch, C., Levine, S., and Hausman, K. Relay policy learning: Solving long-horizon tasks via imitation and reinforcement learning. *arXiv preprint arXiv:1910.11956*, 2019.
- Haarnoja, T., Tang, H., Abbeel, P., and Levine, S. Reinforcement learning with deep energy-based policies. In *International Conference on Machine Learning*, 2017.
- Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning*, 2018a.
- Haarnoja, T., Zhou, A., Hartikainen, K., Tucker, G., Ha, S., Tan, J., Kumar, V., Zhu, H., Gupta, A., Abbeel, P., et al. Soft actor-critic algorithms and applications. Technical report, 2018b.
- Heess, N., Wayne, G., Silver, D., Lillicrap, T. P., Erez, T., and Tassa, Y. Learning continuous control policies by stochastic value gradients. In *Advances in Neural Information Processing Systems*, 2015.
- Jurgenson, T., Avner, O., Groshev, E., and Tamar, A. Sub-goal trees a framework for goal-based reinforcement learning. In *International Conference on Machine Learning*, 2020.
- Kaelbling, L. P. Learning to achieve goals. In *IJCAI*, 1993.
- Kappen, H. J., Gómez, V., and Opper, M. Optimal control as a graphical model inference problem. *Machine learning*, 87(2):159–182, 2012.
- Kostrikov, I., Yarats, D., and Fergus, R. Image augmentation is all you need: Regularizing deep reinforcement learning from pixels. *arXiv preprint arXiv:2004.13649*, 2020.
- Kumar, A., Fu, J., Soh, M., Tucker, G., and Levine, S. Stabilizing off-policy q-learning via bootstrapping error reduction. In *Advances in Neural Information Processing Systems*, 2019.
- Laskin, M., Lee, K., Stooke, A., Pinto, L., Abbeel, P., and Srinivas, A. Reinforcement learning with augmented data. In *Advances in Neural Information Processing Systems*, 2020.
- Levine, S. Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909*, 2018.
- Levine, S. and Abbeel, P. Learning neural network policies with guided policy search under unknown dynamics. In *Advances in Neural Information Processing Systems*, 2014.
- Levine, S. and Koltun, V. Guided policy search. In *International Conference on Machine Learning*, 2013.
- Levine, S., Finn, C., Darrell, T., and Abbeel, P. End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research*, 17(1):1334–1373, 2016.
- Levy, A., Konidaris, G. D., Jr., R. P., and Saenko, K. Learning multi-level hierarchies with hindsight. In *International Conference on Learning Representations*, 2019.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T. P., Harley, T., Silver, D., and Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, 2016.

- Nachum, O., Gu, S., Lee, H., and Levine, S. Data-efficient hierarchical reinforcement learning. In *Advances in Neural Information Processing Systems*, 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*, 2018.
- Nair, A., Dalal, M., Gupta, A., and Levine, S. Accelerating online reinforcement learning with offline datasets. *arXiv preprint arXiv:2006.09359*, 2020.
- Nair, S. and Finn, C. Hierarchical foresight: Self-supervised learning of long-horizon tasks via visual subgoal generation. In *International Conference on Learning Representations*, 2020.
- Nasiriany, S., Pong, V., Lin, S., and Levine, S. Planning with goal-conditioned policies. In *Advances in Neural Information Processing Systems*, 2019.
- Parascandolo, G., Buesing, L., Merel, J., Hasenclever, L., Aslanides, J., Hamrick, J. B., Heess, N., Neitz, A., and Weber, T. Divide-and-conquer monte carlo tree search for goal-directed planning. *arXiv preprint arXiv:2004.11410*, 2020.
- Pertsch, K., Lee, Y., and Lim, J. J. Accelerating reinforcement learning with learned skill priors. *arXiv preprint arXiv:2010.11944*, 2020a.
- Pertsch, K., Rybkin, O., Ebert, F., Zhou, S., Jayaraman, D., Finn, C., and Levine, S. Long-horizon visual planning with goal-conditioned hierarchical predictors. *Advances in Neural Information Processing Systems*, 2020b.
- Peters, J., Mülling, K., and Altun, Y. Relative entropy policy search. In *AAAI Conference on Artificial Intelligence*, 2010.
- Pitis, S., Chan, H., Zhao, S., Stadie, B. C., and Ba, J. Maximum entropy gain exploration for long horizon multi-goal reinforcement learning. In *International Conference on Machine Learning*, 2020.
- 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*, 2018.
- Rawlik, K., Toussaint, M., and Vijayakumar, S. On stochastic optimal control and reinforcement learning by approximate inference. *Proceedings of Robotics: Science and Systems VIII*, 2012.
- Schaul, T., Horgan, D., Gregor, K., and Silver, D. Universal value function approximators. In *International Conference on Machine Learning*, 2015.
- Siegel, N. Y., Springenberg, J. T., Berkenkamp, F., Abdolmaleki, A., Neunert, M., Lampe, T., Hafner, R., and Riedmiller, M. Keep doing what worked: Behavioral modelling priors for offline reinforcement learning. *arXiv preprint arXiv:2002.08396*, 2020.
- Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D., and Riedmiller, M. A. Deterministic policy gradient algorithms. In *International Conference on Machine Learning*, 2014.
- 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.
- Teh, Y. W., Bapst, V., Czarnecki, W. M., Quan, J., Kirkpatrick, J., Hadsell, R., Heess, N., and Pascanu, R. Distral: Robust multitask reinforcement learning. In *Advances in Neural Information Processing Systems*, 2017.
- Tirumala, D., Noh, H., Galashov, A., Hasenclever, L., Ahuja, A., Wayne, G., Pascanu, R., Teh, Y. W., and Heess, N. Exploiting hierarchy for learning and transfer in kl-regularized rl. *arXiv preprint arXiv:1903.07438*, 2019.
- Toussaint, M. Robot trajectory optimization using approximate inference. In *International Conference on Machine Learning*, 2009.
- Veeriah, V., Oh, J., and Singh, S. Many-goals reinforcement learning. *arXiv preprint arXiv:1806.09605*, 2018.
- Vezhnevets, A. S., Osindero, S., Schaul, T., Heess, N., Jaderberg, M., Silver, D., and Kavukcuoglu, K. Feudal networks for hierarchical reinforcement learning. In *International Conference on Machine Learning*, 2017.
- Wang, Z., Novikov, A., Zolna, K., Merel, J., Springenberg, J. T., Reed, S. E., Shahriari, B., Siegel, N. Y., Gülçehre, Ç., Heess, N., and de Freitas, N. Critic regularized regression. In *Advances in Neural Information Processing Systems*, 2020.
- Wiering, M. and Schmidhuber, J. Hq-learning. *Adaptive Behavior*, 6(2):219–246, 1997.
- Wu, Y., Tucker, G., and Nachum, O. Behavior regularized offline reinforcement learning. *arXiv preprint arXiv:1911.11361*, 2019.
- Zhang, L., Yang, G., and Stadie, B. C. World model as a graph: Learning latent landmarks for planning. *arXiv preprint arXiv:2011.12491*, 2020.
- Zhao, R., Sun, X., and Tresp, V. Maximum entropy-regularized multi-goal reinforcement learning. In *International Conference on Machine Learning*, 2019.