
Bayesian Nonparametrics for Offline Skill Discovery

Valentin Villecroze^{1,2} Harry J. Braviner¹ Panteha Naderian¹ Chris J. Maddison^{2,3} Gabriel Loaiza-Ganem¹

Abstract

Skills or low-level policies in reinforcement learning are temporally extended actions that can speed up learning and enable complex behaviours. Recent work in offline reinforcement learning and imitation learning has proposed several techniques for skill discovery from a set of expert trajectories. While these methods are promising, the number K of skills to discover is always a fixed hyperparameter, which requires either prior knowledge about the environment or an additional parameter search to tune it. We first propose a method for offline learning of options (a particular skill framework) exploiting advances in variational inference and continuous relaxations. We then highlight an unexplored connection between Bayesian nonparametrics and offline skill discovery, and show how to obtain a nonparametric version of our model. This version is tractable thanks to a carefully structured approximate posterior with a dynamically-changing number of options, removing the need to specify K . We also show how our nonparametric extension can be applied in other skill frameworks, and empirically demonstrate that our method can outperform state-of-the-art offline skill learning algorithms across a variety of environments. Our code is available at <https://github.com/layer6ai-labs/BNPO>.

1. Introduction

Hierarchical policies have been studied in reinforcement learning for decades (Dietterich et al., 1998; Vezhnevets et al., 2017). Low-level policies in hierarchical frameworks

can be interpreted as skills endowing agents with temporal abstraction. These skills also offer a divide-and-conquer approach to complex reinforcement learning environments, and learning them is therefore a relevant problem.

One of the most commonly used skills frameworks is options (Sutton et al., 1999). Methods for online learning of options have been proposed (Bacon et al., 2017; Khetarpal et al., 2020), but cannot leverage data from expert trajectories. To the best of our knowledge, the only method allowing offline option discovery is that of Fox et al. (2017) (*deep discovery of options*, DDO), which forgoes the use of highly successful variational inference advances (Kingma & Welling, 2013; Rezende et al., 2014) for discrete latent variables (Maddison et al., 2016; Jang et al., 2016) which in this case correspond to the options to be inferred. Our first contribution is proposing a method for offline learning of options combining these previously neglected advances along with a judiciously constructed approximate posterior, which we show empirically outperforms not only DDO, but also other offline skill discovery algorithms.

Additionally, all (discrete) skill learning approaches we are aware of – including options – require specifying the number of skills K as a hyperparameter, rather than learning it. This is a significant practical limitation, as K cannot realistically be known in advance for complex environments. Our second contribution is highlighting a similarity between mixture models and inferring skills from expert trajectories: both are a form of soft clustering where observed data (state-action pairs) are probabilistically assigned to their corresponding cluster (skill). Bayesian nonparametric approaches such as Dirichlet process mixture models (Neal, 2000) are not only mathematically elegant, but also circumvent the practical need to prespecify the number of clusters. We therefore argue that Bayesian nonparametrics should be more heavily used for option and skill discovery.

Our third contribution is proposing a scheme allowing for our option discovery method to accommodate a nonparametric prior over options (thus modelling K as infinite), which can also be applied in other skill learning frameworks. By further adding structure to the variational posterior, allowing it to dynamically change the number of options/skills being assigned positive probability, we recover a method that is not only tractable, but also retains the nonparametric aspect

¹Layer 6 AI, Toronto, Canada ²University of Toronto, Toronto, Canada ³Vector Institute, Toronto, Canada. Correspondence to: Valentin Villecroze <valentin.v@layer6.ai>, Harry J. Braviner <harry@layer6.ai>, Panteha Naderian <panteha@layer6.ai>, Chris J. Maddison <cmaddis@cs.toronto.edu>, Gabriel Loaiza-Ganem <gabriel@layer6.ai>.

of the model and does not require treating K as a hyperparameter. We show that our nonparametric versions of skill learning algorithms match the performance of their parametric counterparts with tuned K , thus successfully “learning a sufficient number of options” in practice. Finally, we hope that our nonparametric variational inference scheme will find uses outside of offline skill recovery.

2. Related Work

Imitation learning (Hussein et al., 2017), also called *learning from demonstrations* (Schaal et al., 1997; Argall et al., 2009), consists of learning to solve a task from a set of expert demonstrations. This can be achieved by methods such as behavioural cloning (Esmaili et al., 1995; Ross et al., 2011) or inverse reinforcement learning (Ho & Ermon, 2016; Sharma et al., 2018; Krishnan et al., 2019). In this paper we focus on the former. Many recent works on imitation learning attempt to divide the demonstrations into several smaller segments, before fitting models on each of them (Niekum et al., 2013; Murali et al., 2016; Krishnan et al., 2018; Shiarlis et al., 2018). These works usually consider these two steps as two distinct stages, and unlike our proposed approach, do not learn skills end-to-end. Niekum et al. (2013) and Krishnan et al. (2018) are of particular interest as they are the only ones, to our knowledge, to make use of a Bayesian nonparametric model to segment the trajectories. However, they significantly differ from our method in that they use handcrafted movement primitives and linear dynamical models to fit the resulting segments.

Our work is closer to the field of option and skill discovery, which leverages work on online hierarchical reinforcement learning (Sutton et al., 1999; Gregor et al., 2016; Bacon et al., 2017; Achiam et al., 2018; Eysenbach et al., 2018; Sharma et al., 2019; Florensa et al., 2017) and aims to learn adequate options from a set of expert trajectories. DDO (Fox et al., 2017) uses an EM-type algorithm (Dempster et al., 1977) that allows multi-level hierarchies. As previously mentioned, this method forgoes advances in variational inference and continuous relaxations, and we will later show that option learning can be significantly improved upon through the use of these techniques. While not following the option framework, Kipf et al. (2019) propose CompILE, which uses a variational approach and continuous relaxations to both segment the trajectories and encode the resulting slices as discrete skills. We will show that with these advances, options also outperform CompILE.

Both DDO and CompILE need the number of options/skills, K , to be specified beforehand. We borrow from the Bayesian nonparametrics literature (Teh et al., 2006; Caron et al., 2007; Dunson & Xing, 2009) – in particular from models involving Dirichlet processes (Ferguson, 1973) – in order to model options. We follow Nalisnick & Smyth

(2016) and use a variational-autoencoder-inspired (Kingma & Welling, 2013; Rezende et al., 2014) inference scheme that avoids the need for the expensive MCMC samplers commonly associated with these types of models (Neal, 2000). As a result, the nonparametric versions of the models we consider do not need to prespecify nor tune K .

Finally, Shankar & Gupta (2020) and Ajay et al. (2021) also use variational inference to discover skills but choose to represent them as continuous latent variables instead of categorical variables. While this choice also technically leads to infinitely many skills, as we will see in our experiments, discrete skills learned from expert trajectories are a useful way to enhance online agents. The use of continuous skills results not only in less interpretable skills, but also loses the ability to perform these enhancements.

3. Preliminaries

We now review the main concepts needed for our model: behavioural cloning, the options framework, continuous relaxations, and the stick-breaking process. We also review CompILE, a method that we shall later make nonparametric.

3.1. Behavioural Cloning

The goal of behavioural cloning is to imitate an expert who demonstrates how to accomplish a task. More formally, we consider a Markov Decision Process without rewards (MDP\R). An MDP\R is a tuple $\mathcal{M} : \langle \mathcal{S}, \mathcal{A}, P, \rho \rangle$, where \mathcal{S} is the state space, \mathcal{A} the action space, $P(s_{t+1} | s_t, a_t)$ a transition distribution, and ρ the starting state distribution. We assume access to a dataset of expert trajectories $\xi := \{\xi^{(i)} := (s_0^{(i)}, a_0^{(i)}, s_1^{(i)}, a_1^{(i)}, \dots, s_T^{(i)})\}_{i=1}^N$ of states $s_t^{(i)} \in \mathcal{S}$ and actions $a_t^{(i)} \in \mathcal{A}$. The trajectory length T need not be identical across trajectories and could be changed to T_i , but we keep T to avoid further complicating notation. We want to find a policy $\pi_\theta : \mathcal{S} \mapsto \Delta(\mathcal{A})$ (where $\Delta(\mathcal{A})$ denotes the set of distributions over \mathcal{A}) parameterized by θ that maximizes the logarithm of following expression:

$$p_\theta(\xi) = \prod_{i=1}^N \rho(s_0^{(i)}) \prod_{t=0}^{T-1} \pi_\theta(a_t^{(i)} | s_t^{(i)}) P(s_{t+1}^{(i)} | s_t^{(i)}, a_t^{(i)}) \quad (1)$$

$$\propto \prod_{i=1}^N \prod_{t=0}^{T-1} \pi_\theta(a_t^{(i)} | s_t^{(i)}), \quad (2)$$

i.e. the policy maximizing the likelihood assigned to the expert trajectories. We have dropped terms that are constant with respect to θ in Equation 2.

3.2. The Options Framework

As previously mentioned, the option framework introduces temporally extended actions, allowing for temporal abstraction and more complex behaviour. Formally, an option

Algorithm 1 Trajectory generation with options.

```

1:  $s_0 \sim \rho(\cdot), b_0 \leftarrow 1$ 
2: for  $t \in [0, \dots, T - 1]$  do
3:   if  $b_t = 1$  then
4:      $h_t \sim \eta(\cdot|s_t) \setminus \setminus \text{Draw option from high-level policy}$ 
5:   else
6:      $h_t \leftarrow h_{t-1}$ 
7:   end if
8:    $a_t \sim \pi_{h_t}(\cdot|s_t) \setminus \setminus \text{Draw action from low-level policy}$ 
9:    $s_{t+1} \sim P(\cdot|s_t, a_t) \setminus \setminus \text{Draw the next state}$ 
10:   $b_{t+1} \sim \text{Bernoulli}(\cdot|\psi_{h_t}(s_{t+1})) \setminus \setminus \text{Terminate option?}$ 
11: end for
    
```

$h \in \Omega$ (Sutton et al., 1999) is a tuple $\langle \mathcal{I}_h, \pi_h, \psi_h \rangle$, where $\mathcal{I}_h \subseteq \mathcal{S}$ is the initiation set, $\pi_h(a|s)$ the control or low-level policy, and $\psi_h : \mathcal{S} \rightarrow [0, 1]$ the termination function. An option h can be invoked in any state $s \in \mathcal{I}_h$, in which case actions are drawn according to π_h until the option terminates, which happens with probability $\psi_h(s^0)$ at each subsequent state s^0 . In our setting we assume that $\mathcal{I}_h = \mathcal{S}$, which is a common assumption (Bacon et al., 2017; Fox et al., 2017; Shankar & Gupta, 2020). In order to use these options to solve a task, a high-level policy (or policy over options) $\eta : \mathcal{S} \mapsto \Delta(\Omega)$ is used to select a new option h after the previous one terminates. This two-level structure is called a hierarchical policy. The resulting generative process is described in Algorithm 1.

Learning a hierarchical policy from expert demonstrations using behavioural cloning is much more challenging than with a single policy. Indeed, not only do multiple policies need to be learned concurrently (one for each option), but, as the options are unobserved, they must be treated as latent variables. Fox et al. (2017) propose an EM-based approach to do this, which we improve upon.

3.3. CompILE

CompILE (Kipf et al., 2019) is used for hierarchical imitation learning but does not rely on the option framework. Binary termination variables are not sampled at every timestep to determine if the current skill should terminate. Rather, whenever a new skill is drawn from the high-level policy, CompILE also samples the number of steps for which the skill will be active from a Poisson distribution with parameter λ . Also, rather than using a state-dependent high-level policy, η is assumed to be uniform over skills. We refer to any time interval between subsequent skill selections as a segment. This procedure is summarized in Algorithm 2, where we have abused notation and still use b to denote termination variables in order to highlight the similarity with the options framework, even though these are not the same variables as in options. Similarly, h_j does not denote an option, but rather the skill being used during segment j .

Algorithm 2 Trajectory generation with CompILE.

```

1:  $s_0 \sim \rho(\cdot), b_0 \leftarrow 0, j \leftarrow 0$ 
2: for  $t \in [0, \dots, T - 1]$  do
3:   while  $t = b_j$  do
4:      $j \leftarrow j + 1 \setminus \setminus \text{Consider the next segment}$ 
5:      $h_j \sim \eta(\cdot) = \text{Uniform}(\cdot|K) \setminus \setminus \text{Draw skill}$ 
6:      $b_j \sim \text{Poisson}(\cdot|\lambda) \setminus \setminus \text{Draw segment length}$ 
7:      $b_j \leftarrow b_j + b_{j-1} \setminus \setminus \text{Next segment boundary}$ 
8:   end while
9:    $a_t \sim \pi_{h_j}(\cdot|s_t) \setminus \setminus \text{Draw action from low-level policy}$ 
10:   $s_{t+1} \sim P(\cdot|s_t, a_t) \setminus \setminus \text{Draw the next state}$ 
11: end for
    
```

Doing behavioural cloning in CompILE is harder than just maximizing Equation 2, since computing $\log p_\theta(\xi)$, where θ now parameterizes all the low-level policies, requires an intractable marginalization over the unobserved variables (b 's and h 's). In order to circumvent this issue, an approximate posterior $q_\phi(\zeta|\xi)$ is introduced, where ζ denotes all the unobserved variables for all trajectories and ϕ the variational parameters. Instead of directly maximizing $\log p_\theta(\xi)$, the following lower bound, which is called the ELBO, is maximized over (θ, ϕ) :

$$\begin{aligned}
 L(\theta, \phi) &:= \mathbb{E}_{q_\phi(\zeta|\xi)}[\log p_\theta(\zeta, \xi) - \log q_\phi(\zeta|\xi)] \quad (3) \\
 &\leq \log p_\theta(\xi). \quad (4)
 \end{aligned}$$

We omit the details on how $q_\phi(\zeta|\xi)$ is structured, but highlight that the number of segments needs to be specified as a hyperparameter of this variational approximation rather than being properly treated as random. In order to maximize Equation 3, the authors use the reparameterization trick of variational autoencoders (Kingma & Welling, 2013; Rezende et al., 2014) along with continuous relaxations. In particular, they use the Concrete or Gumbel-Softmax (GS) distribution (Maddison et al., 2016; Jang et al., 2016), which we briefly review in the next section, in order to efficiently backpropagate through Equation 3.

3.4. Continuous Relaxations

The reparameterization trick is used in variational autoencoders to backpropagate through expressions of the form $\mathbb{E}_{q_\phi(\zeta)}[f_\phi(\zeta)]$, with respect to ϕ , where f_ϕ is a real-valued function. Since the distribution with respect to which the expectation is taken depends on ϕ , one cannot simply bring the gradient inside of the expectation. Gradient estimators such as REINFORCE (Glynn, 1990; Williams, 1992) typically exhibit high variance, which the reparameterization trick empirically reduces. When ζ is a continuous random variable, it is often (but not always) the case that one can easily find a continuously differentiable function g such that $\zeta \sim q_\phi(\cdot) \iff \zeta = g(\epsilon, \phi)$ where ϵ follows some continuous distribution which does not depend on ϕ . In this case,

the gradient is given by:

$$\nabla_{\phi} \mathbb{E}_{q_{\phi}(\zeta)}[f_{\phi}(\zeta)] = \mathbb{E}_{\epsilon}[\nabla_{\phi} f_{\phi}(g(\epsilon, \phi))], \quad (5)$$

and a Monte Carlo estimate can be easily obtained.

When ζ is categorical, g has to be piece-wise constant and Equation 5 no longer holds. Continuous relaxations approximate $\mathbb{E}_{q_{\phi}(\zeta)}[f_{\phi}(\zeta)]$ with an expectation over a continuous random variable, so that the reparameterization trick can be used. The Gumbel-Softmax (GS) or Concrete distribution is a distribution on the K -simplex, $\Delta(K) := \{x \in \mathbb{R}^K : x_k > 0, \sum_{k=0}^K x_k = 1\}$, parameterized by $q \in \Delta(K)$ and a temperature hyperparameter $\tau > 0$, designed to address this issue. It is reparameterized as follows:

$$\zeta \sim GS_{\tau}(\cdot|q) \iff \zeta = \text{softmax}\left(\frac{\epsilon + \log q}{\tau}\right), \quad (6)$$

where ϵ is a K -dimensional vector with independent Gumbel(0, 1) entries, and the log is taken elementwise. As $\tau \rightarrow 0$, the $GS_{\tau}(\cdot|q)$ distribution converges to the discrete distribution q . By thinking of categorical ζ 's as one-hot vectors of length K , the GS thus provides a continuous relaxation of ζ , and $\mathbb{E}_{q_{\phi}(\zeta)}[f_{\phi}(\zeta)]$ can be approximated by:

$$\mathbb{E}_{q_{\phi}(\zeta)}[f_{\phi}(\zeta)] \approx \mathbb{E}_{GS_{\tau}(\zeta|q_{\phi}(\cdot))}[\tilde{f}(\zeta)], \quad (7)$$

where we think of the discrete distribution $q_{\phi}(\cdot)$ as a K -dimensional vector, and \tilde{f} is a relaxation of f mapping all of $\Delta(K)$ (rather than just the vertices, i.e. one-hot vectors) to \mathbb{R} . Equation 7 admits the gradient estimator of Equation 5, and Kipf et al. (2019) use it to optimize Equation 3.

3.5. The Stick-Breaking Process

The stick-breaking process (Sethuraman, 1994; Ishwaran & James, 2001), which is deeply connected to Dirichlet processes, places a distribution over probability vectors of infinite length. Equivalently, it is a distribution over distributions on a countably infinite set. This process will allow us to both assume that there are infinitely many options, and to place a proper prior on the high-level policy η itself (i.e. a distribution over distributions of options). When performing posterior inference some options will have extremely small probabilities; this enables us to “learn the number of options” from the data. More formally, the Griffiths-Engen-McCloskey distribution, denoted GEM(α) and parameterized by $\alpha > 0$, produces an infinitely long vector $(\beta_0, \beta_1, \dots)$ such that $\beta_k > 0$ and $\sum_{k=0}^{\infty} \beta_k = 1$ almost surely. Samples are obtained by first sampling $\beta_k^0 \sim \text{Beta}(1, \alpha)$ for $k = 0, 1, \dots$ and then setting:

$$\beta_k = \beta_k^0 \prod_{l=0}^{k-1} (1 - \beta_l^0). \quad (8)$$

Intuitively, we start off with a stick of length 1, and at every step we remove β_k from the stick. β_k^0 is the fraction of the remaining stick that is assigned to β_k .

We shall later use the GEM distribution as a prior, and will want to perform variational inference. Applying the stick-breaking procedure to Beta random variables with different parameters might seem like the most natural approximate posterior. However, the Beta distribution is not easily reparameterized, and thus inference with such a posterior becomes computationally challenging. It is therefore common to instead use the Kumaraswamy distribution in the approximate posterior (Nalisnick & Smyth, 2016; Stirn et al., 2019). Like the Beta distribution, this is a (0, 1)-valued distribution with two parameters $a_1, a_2 > 0$, whose density is given by:

$$p(x|a_1, a_2) = a_1 a_2 x^{a_1-1} (1-x)^{a_2-1}, \quad (9)$$

and which can easily be reparameterized.

4. Our Model

We first describe a parametric version of our model, where the number of options K is a fixed hyperparameter. We will detail in Sections 5.1 and 5.2 how we use a nonparametric prior to circumvent the need to specify this number. Throughout this section and until Section 5.3 (exclusive), our notation refers to options and not to CompILE.

4.1. Overview

Now that we have covered the required concepts, we introduce our method in more detail. We assume that the expert trajectories ξ are generated by a two-level option hierarchy, as described in Algorithm 1, with the high-level policy η being shared across trajectories. We denote the corresponding trajectories of hidden options h and binary termination variables b as $\zeta := \{\zeta^{(i)} := (b_0^{(i)}, h_0^{(i)}, b_1^{(i)}, h_1^{(i)}, \dots, h_{T-1}^{(i)})\}_{i=1}^N$, and η the high-level policy. Fox et al. (2017) and Kipf et al. (2019) found that taking η as a uniform policy rather than learning it results in equally useful learned skills. We simplify η with the less restrictive assumption that it does not depend on the current state s , a choice that will later simplify our nonparametric component. That is, in Algorithm 1, we draw h_t according to $\eta(\cdot)$ rather than $\eta(\cdot|s_t)$. We then treat η in a Bayesian way, as a latent variable to be inferred from observed trajectories, and assume a K -dimensional prior over it, obtained by truncating the stick-breaking process of Beta(1, α) variables after $K - 1$ steps (and having the last entry be such that the resulting vector adds up to 1). The resulting graphical model can be seen in Figure 1.

We denote as θ all the parameters from the prior (i.e. α), the low-level policies, and the terminations functions. We implement the termination functions with a single neural

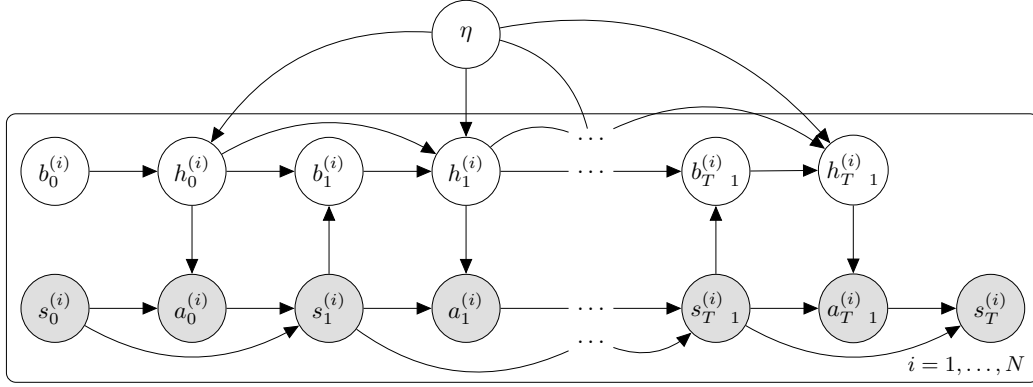


Figure 1. Graphical model for options. Shaded nodes correspond to observed variables ($\xi^{(i)}$) while the blank ones are latent ($\zeta^{(i)}$).

network taking s as input and outputting the termination probabilities for every option. We keep a separate neural network for every low-level policy, each of which takes s as input and outputs the parameters of a categorical over actions (we assume a discrete action space, but our method can easily be modified for continuous action spaces). Note that, other than the Bayesian treatment of η and the inclusion of α in θ , our model is identical to DDO. We will detail differences in how learning is performed in this section, and later show how our model can be made nonparametric.

As with CompILE, naïvely doing behavioural cloning is intractable, and we therefore introduce an approximate posterior $q_\phi(\eta, \zeta|\xi)$. As we will cover in detail later, we endow this posterior with a structure respecting the conditional independences implied by Figure 1, and unlike CompILE, we do not have to arbitrarily specify a number of segments.

4.2. Objective

We jointly train θ and ϕ by using the ELBO as the maximization objective, which is now given by:

$$L(\theta, \phi) := \mathbb{E}_{q_\phi(\eta, \zeta|\xi)} [\log p_\theta(\zeta, \xi|\eta) - \log q_\phi(\zeta|\eta, \xi)] - \text{KL}(q_\phi(\eta)||p(\eta)) \leq \log p_\theta(\xi). \quad (10)$$

We use the GS distribution in order to perform the reparameterization trick and backpropagate through the ELBO. If we write the first term in more detail, we get:

$$\log p_\theta(\zeta, \xi|\eta) = \sum_{i=1}^N \log p_\theta(\zeta^{(i)}, \xi^{(i)}|\eta), \quad (11)$$

where each term is given by:

$$\begin{aligned} \log p_\theta(\zeta^{(i)}, \xi^{(i)}|\eta) &= \log p(s_0^{(i)}) + \log \delta_{b_0^{(i)}=1} \\ &+ \log \eta(h_0^{(i)}) + \sum_{t=1}^{T-1} \log p_\theta(b_t^{(i)}, h_t^{(i)}|h_{t-1}^{(i)}, s_t^{(i)}, \eta) \\ &+ \sum_{t=0}^{T-1} \log \pi_{h_t^{(i)}}(a_t^{(i)}|s_t^{(i)}) + \log P(s_{t+1}^{(i)}|s_t^{(i)}, a_t^{(i)}). \end{aligned} \quad (12)$$

Note that the initial state distribution and environment dynamics terms cannot be evaluated but, as they do not depend on the parameters being optimized, they can be ignored. We can further decompose the term for $b_t^{(i)}$ and $h_t^{(i)}$:

$$p_\theta(b_t^{(i)} = 1, h_t^{(i)}|h_{t-1}^{(i)}, s_t^{(i)}, \eta) = \psi_{h_{t-1}^{(i)}}(s_t^{(i)})\eta(h_t^{(i)}) \quad (13)$$

$$\begin{aligned} p_\theta(b_t^{(i)} = 0, h_t^{(i)}|h_{t-1}^{(i)}, s_t^{(i)}, \eta) &= (1 - \psi_{h_{t-1}^{(i)}}(s_t^{(i)})) \\ &\cdot \delta_{h_t^{(i)}=h_{t-1}^{(i)}}, \end{aligned} \quad (14)$$

where $\{\psi_h\}_h$ are the termination functions. Note that equations 12, 13 and 14 assume that the b 's are binary and the h 's are categorical (one-hot) through the Kronecker delta terms and evaluating η at different h terms. Since we use a Gumbel-Softmax relaxation to optimize the ELBO, the sampled values of b 's and h 's will be $(0, 1)$ and simplex-valued, respectively. We therefore relax the corresponding terms so that they can be evaluated for such values and gradients can be propagated. We detail the relaxations in Appendix A.

4.3. Variational Posterior

We now describe the structure of our approximate posterior $q_\phi(\eta, \zeta|\xi)$. First, we observe from Figure 1 through the rules of d -separation (Koller & Friedman, 2009) that $\zeta^{(i)}$ is independent of $\zeta^{(j)}$ given η and ξ , provided $i \neq j$. Additionally, given η , $\zeta^{(i)}$ depends on ξ only through $\xi^{(i)}$. We thus take $q_\phi(\eta, \zeta|\xi)$ to obey the following conditional independence relationship:

$$q_\phi(\eta, \zeta|\xi) = q_\phi(\eta) \prod_{i=1}^N q_\phi(\zeta^{(i)}|\eta, \xi^{(i)}), \quad (15)$$

which holds for the true posterior as well.

Since η is a global variable (i.e. it does not have components for each trajectory i) while $\zeta = \{\zeta^{(i)}\}_i$ is composed of local variables assigned to the corresponding $\xi^{(i)}$, we treat η in a non-amortized (Gershman & Goodman, 2014) way (i.e.

the parameters of $q_\phi(\eta)$ are directly optimized, there is no neural network), while we treat $\zeta^{(i)}$ in an amortized way (i.e. a neural network takes η and $\xi^{(i)}$ as inputs to produce the parameters of a distribution).

As previously hinted, we parameterize $q_\phi(\eta)$ as a sequence of Kumaraswamy distributions to which the stick-breaking procedure is applied, which allows us to straightforwardly use the reparameterization trick.

As mentioned above, $q_\phi(\zeta^{(i)}|\eta, \xi^{(i)})$ is treated in an amortized way, so that a neural network takes η and $\xi^{(i)}$ as input and produces the parameters for the distribution of $\zeta^{(i)}$, i.e. the Bernoulli parameters of each $b_t^{(i)}$ and categorical parameters of $h_t^{(i)}$ for $t = 0, \dots, T-1$. We use an autoregressive structure for $q_\phi(\zeta^{(i)}|\eta, \xi^{(i)})$, and assume conditional independence of options and terminations at every time step:

$$q_\phi(\zeta^{(i)}|\eta, \xi^{(i)}) = q_\phi(b_0^{(i)}|\eta, \xi^{(i)})q_\phi(h_0^{(i)}|\eta, \xi^{(i)}) \prod_{t=1}^{T-1} q_\phi(b_t^{(i)}|\zeta_{t-1}^{(i)}, \eta, \xi_{t:T-1}^{(i)})q_\phi(h_t^{(i)}|\zeta_{t-1}^{(i)}, \eta, \xi_{t:T-1}^{(i)}), \quad (16)$$

where $\xi_t^{(i)} = (s_t^{(i)}, a_t^{(i)})$ and $\zeta_t^{(i)} = (b_t^{(i)}, h_t^{(i)})$. Note that conditioned on $(\zeta_{t-1}^{(i)}, \eta, \xi^{(i)})$, $(b_t^{(i)}, h_t^{(i)})$ is independent of $\xi_{0:t-1}$, which can again be verified through Figure 1. In other words, states and actions after time t can be useful to infer the option and termination at time t ; but previous states and actions are not (assuming the previous option and termination are known). This is why we only condition on $\xi_{t:T-1}^{(i)}$ in Equation 16, rather than on $\xi^{(i)}$. In practice, we use an LSTM (Hochreiter & Schmidhuber, 1997) to parse $\xi^{(i)}$ in reverse order, then sequentially sample $\zeta^{(i)}$ using multi-layer perceptrons (MLPs) to output the distributions' parameters as shown below (the MLPs' inputs are concatenated):

$$b_t^{(i)}|\zeta_{t-1}^{(i)}, \eta, \xi^{(i)} \sim \text{GS}_\tau(\text{MLP}_b(\text{LSTM}(\xi_{t:T-1}^{(i)}), \eta, \zeta_{t-1}^{(i)})), \quad (17)$$

$$h_t^{(i)}|\zeta_{t-1}^{(i)}, \eta, \xi^{(i)} \sim \text{GS}_\tau(\text{MLP}_h(\text{LSTM}(\xi_{t:T-1}^{(i)}), \eta, \zeta_{t-1}^{(i)})). \quad (18)$$

The b and h indices reflect the fact that we use two separate output heads. The MLPs share their parameters except those of their last linear layer. The $\text{LSTM}(\xi_{t:T-1}^{(i)})$ term denotes the hidden state of the LSTM layer taken at time step t . Also note that we abuse notation and use $b_t^{(i)}$, $h_t^{(i)}$ and $\zeta_t^{(i)}$ for both the discrete variables and their relaxed counterparts. We refer to the LSTM and the MLP heads as the *encoder*.

4.4. Entropy Regularizer

In some preliminary experiments, we observed that our model could get stuck in some local optima where too few options were used. To address this issue, we add a regularizing term to the ELBO in Equation 10, defined as the entropy over the average of the sampled options $\{h_t^{(i)}\}$. More formally, we define $h_{avg}^{(i)} := \frac{1}{T} \sum_{t=0}^{T-1} h_t^{(i)}$, which is

a probability vector of size K and consider the following regularizing term that we want to maximize:

$$l_{ent} := -\mathbb{E}_{q_\phi(\zeta|\xi)} \left[\sum_{i=1}^N \sum_{k=0}^{K-1} h_{avg,k}^{(i)} \log \left(h_{avg,k}^{(i)} \right) \right], \quad (19)$$

where $h_{avg,k}^{(i)}$ is the k -th coordinate of $h_{avg}^{(i)}$. We also use the reparameterization trick to backpropagate through Equation 19. This term encourages the model to use all the available options in equal amounts. Our final maximization objective is given by the following expression:

$$L(\theta, \phi) + \lambda_{ent} l_{ent}, \quad (20)$$

where we anneal $\lambda_{ent} \geq 0$ throughout training with a fixed decay rate. We highlight that we only anneal λ_{ent} since the ELBO is a highly principled objective, and annealing causes the objective being optimized to be closer and closer to the ELBO. That being said, we did not observe significant empirical differences between annealing and fixing λ_{ent} throughout training (see Appendix D).

5. Nonparametric Models

We now describe how to use a nonparametric prior to remove the need for a prespecified number of options.

5.1. GEM Prior

We now put a GEM(α) prior over η , replacing the prior obtained from truncating the stick-breaking process at K steps. The result is a nonparametric model which assumes a countably infinite number of options. Doing posterior inference over η then enables learning the number of options present in the observed trajectories.

Note that θ should have an infinite number of parameters, since we consider an infinite number of options, but we detail below how we manage this in practice.

5.2. Truncation

Recall that η is now an infinite dimensional vector whose entries add up to one. A naïve attempt to optimize the objective from Equation 10 would thus involve not only sampling infinitely long η 's, but also neural networks with infinitely many parameters (for terminations), or infinitely many neural networks (for low-level policies); this is clearly infeasible. Nalisnick & Smyth (2016) truncate η and consider only the K first coordinates, treating K as a hyperparameter. We highlight that fixing K in the approximate posterior yields the same objective as reverting the nonparametric GEM prior back to a K -dimensional prior; thus discarding the nonparametric aspect of the model. Even in this setting, our fixed- K model remains a novel way of learning options offline. Nonetheless, in order to properly retain

the nonparametric aspect of the model, we allow K to increase throughout training. We check at regular intervals (every n_K training steps) whether we should increment K by looking at the “usage”, $U(h)$, for each option h :

$$U(h) := \frac{1}{NT} \sum_{i=1}^N \sum_{t=0}^{T-1} \mathbb{1} \left(h = \underset{h_k: 0 \leq k < K}{\operatorname{argmax}} \pi_{h_k}(a_t^{(i)} | s_t^{(i)}) \right), \quad (21)$$

i.e. $U(h)$ is the fraction of steps for which we inferred h to be the option most likely to have been in use. Our rule for adding skills (clusters) is inspired by the small variance asymptotics Bayesian nonparametrics literature (Kulis & Jordan, 2011; Broderick et al., 2013). We increase K by 1 if there is no option h_k for $k \in \{0, \dots, K-1\}$ such that $U(h_k) < \delta/K$ for a fixed hyperparameter $\delta < 1$. When increasing K , we add a new low-level policy, and expand the termination function by adding a new row to the last linear layer, so that the output changes from size K to size $K+1$. We handle this new option in the encoder’s MLPs (see Equations 17 and 18) by adding two new columns to their first layer (to take into account the new input size, as both η and $\zeta_t^{(i)}_1$ have increased in length by 1) and a new row to the last layer of MLP $_h$ (to output the correct number of parameters for $\zeta_t^{(i)}$). We also add parameters for another Kumaraswamy distribution to ϕ . All new parameters are initialized randomly. The logic behind this heuristic is that there is no need to add a new option (i.e. increment K) if there is an existing option which is rarely used.

5.3. Nonparametric CompILE

Finally, we show that CompILE can now be easily turned into a nonparametric model thanks to the machinery we have developed. When placing a GEM(α) prior over η , the CompILE ELBO from Equation 3 changes to accommodate the fact that η is now being treated in a Bayesian way rather than being fixed and uniform. The ELBO then becomes:

$$L(\theta, \phi) := \mathbb{E}_{q_\phi(\eta, \zeta | \xi)} [\log p_\theta(\zeta, \xi | \eta) - \log q_\phi(\zeta | \eta, \xi)] - \mathbb{KL}(q_\phi(\eta) || p(\eta)). \quad (22)$$

While this objective looks identical to ours, we highlight once again that we have switched back to CompILE notation, and that the terms are not identical to those of Equation 10. The only practical changes to CompILE are then: (1) that we also use a non-amortized posterior $q_\phi(\eta)$ obtained by stick-breaking Kumaraswamy random variables, and so the \mathbb{KL} term in the ELBO is identical to that of Equation 10; (2) the approximate posterior becomes $q_\phi(\eta, \zeta | \xi) = q_\phi(\eta)q_\phi(\zeta | \eta, \xi)$, which implies that, as with our own model, the neural network modeling $q_\phi(\zeta | \eta, \xi)$ now needs to take η as an extra input, although we leave the rest of the structure used in CompILE intact; and (3) when evaluating $\eta(h_j^{(i)})$ for $p_\theta(\zeta, \xi | \eta)$, where $h_j^{(i)}$ is the skill used for the j -th segment in the i -th trajectory, we actually evaluate

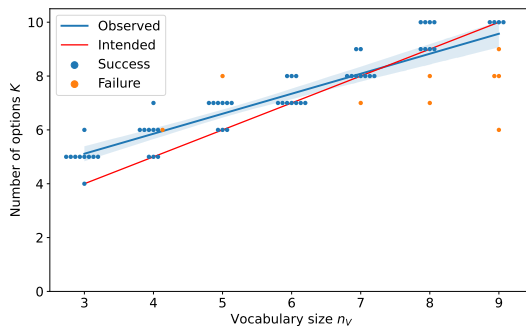


Figure 2. Results for our proof-of-concept environment. We consider 10 random seeds per vocabulary size n_V , and a run is considered a success if one of the learned options selects the correct action at $t = 4$ with probability at least 0.95. The red line shows the number of options that should be recovered ($n_V + 1$), while the blue line is the linear regression of the actual recovered numbers and closely mimics the red line.

η (as we do in options), rather than using $1/K$ as implied by the uniform distribution. Additionally, for the nonparametric version of CompILE we use the same truncation and entropy regularizer as we used for our options model.

6. Experiments

The goal of our experiments is twofold: to show that our options framework learns more useful skills than DDO and CompILE, and also that the nonparametric extensions of our own model and CompILE (which circumvent the need to specify K) match the performance of their respective parametric versions with K tuned as a hyperparameter. The former goal highlights the usefulness of incorporating variational inference advances to offline option learning, and the latter highlights the benefits of using Bayesian nonparametrics for skill discovery. All experimental details are given in Appendix C.

6.1. Proof-of-Concept Environment

We evaluate our model in an environment we designed, where we know the number of skills required to fully reconstruct the expert trajectories. Our nonparametric model recovers sufficient options for reconstruction, without this number being specified beforehand. In this environment, an agent receives a message m from a vocabulary $\{0, \dots, n_V - 1\}$ at time $t = 0$ and has to emit the same message m at time $t = 4$. The agent’s observation is (t, m) at $t = 0$ and $(t, -1)$ for all subsequent timesteps, and they can take a discrete action $a_t \in \{0, \dots, n_V - 1\}$ at each timestep t . The task is considered a success if $a_4 = m$. Note that this task cannot be solved by a Markovian policy,

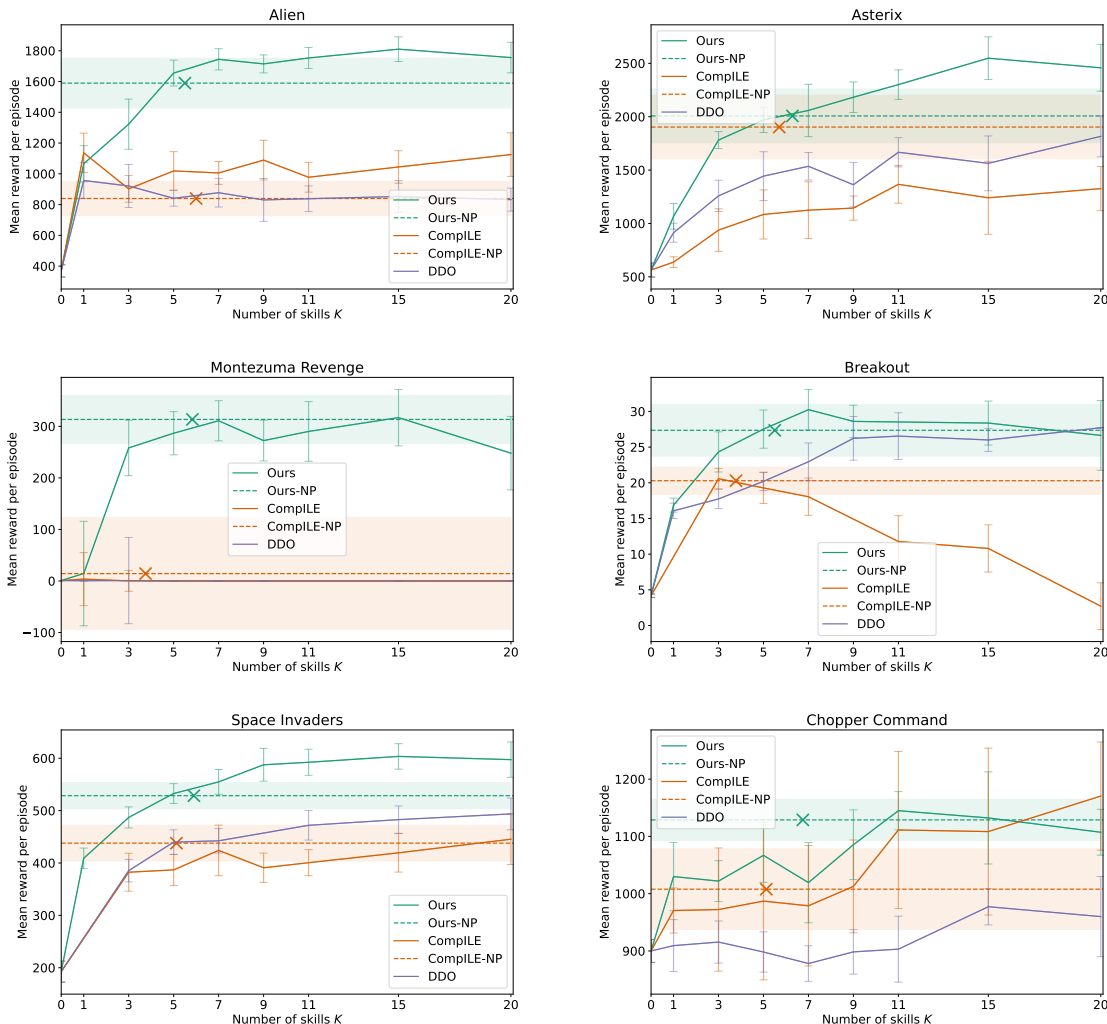


Figure 3. Results for Atari environments showing mean reward per episode (after training) averaged across at least 10 seeds for different numbers of skills, with corresponding standard errors. Dashed lines correspond to nonparametric models where K is not specified. The label ‘ours’ denotes the fixed- K version of our options model, and the ‘-NP’ suffix denotes nonparametric versions, both for ours and CompILE.

since at time $t = 4$ the agent needs to remember the message they received at time $t = 0$. However, it is possible to solve the environment with a hierarchical policy with at least n_V options: if every possible message has a corresponding option whose low-level policy emits that message at time $t = 4$, then an optimal agent needs just select the appropriate option at time $t = 0$ when observing m and not terminate before time $t = 4$. We generate expert trajectories for various vocabulary sizes. Considering the heuristic we use to update K (see Section 5.2), we expect our method to recover at least $K = n_V + 1$ (K options equally used, and one unused option). We show the results in Figure 2: indeed, we recover enough options without overestimating K by more than 1. We also perform ablations in Appendix

D, where we show that removing the entropy regularizer from Section 4.4 significantly hurts performance.

6.2. Atari Environments

We further test our model on several games from the Atari learning environment (Bellemare et al., 2013). For each game, we use expert trajectories generated by a trained Apex agent (Horgan et al., 2018; Such et al., 2019). To evaluate the quality of the discovered options, we train a PPO agent (Schulman et al., 2017) on an environment for which the action space has been ‘augmented’ using the skills, similarly to Fox et al. (2017). In this augmented environment, alongside the default game actions, the agent can choose to perform ‘enhanced actions’, which consist of acting ac-

cording to one of the learned skills until it terminates. We modify the PPO algorithm to take these enhanced actions into account (see Appendix B). We compare our method to DDO and CompILE. We compare both the parametric version of our model where K is treated as a hyperparameter, and the nonparametric versions of our model and CompILE. Results are summarized in Figures 3 and 4. We can see that (1) our model learns more useful skills than DDO and CompILE (green lines consistently above purple and orange ones, except for the Asteroids environment), and that (2) the nonparametric version of our model matches the performance of tuning K (dashed green lines match the peak of the green lines for all environments, except Asterix and Space Invaders), and similarly for CompILE’s nonparametric version (analogously for orange lines, except for Alien and Chopper Command this time). Note that the performance of the agents in the original (i.e. not augmented) versions of the environments correspond to the ‘ $K = 0$ ’ entries in the figures: the significant improvements with $K > 0$ highlight the relevance of offline skill discovery. The crosses in Figures 3 and 4 show the recovered values of K (averaged across runs) for the nonparametric models, and we can see that, except for a few cases, the number of recovered options closely matches the smallest optimal number that would be obtained by tuning K through many runs, underlining the benefits of our nonparametric approaches. We include visualizations of our learned options in the github repository containing our code.¹

We highlight that we did not tune λ_{ent} nor its decay schedule and used the same settings across all experiments, so that not tuning K does not come at the cost of having to tune other parameters. We present ablations with respect to the entropy regularizer in Appendix D, where we found that while it did not hurt performance, this regularizer was not as fundamental as in the proof-of-concept environment. We also include in Appendix D ablations over trajectory length, and find that while results vary for all methods, ours remains the strongest regardless of trajectory length.

7. Conclusions and Future Work

We introduced a novel approach for offline option discovery taking advantage of variational inference and the Gumbel-Softmax distribution. We leave an exploration of continuous relaxation (Kool et al., 2020; Potapczynski et al., 2020; Paulus et al., 2020) and variational inference advances (Burda et al., 2015; Roeder et al., 2017) within our framework for future research.

We also highlighted an unexplored connection between skill discovery and Bayesian nonparametrics, and showed that

¹<https://github.com/layer6ai-labs/BNPO/blob/main/visualization/visualization.md>

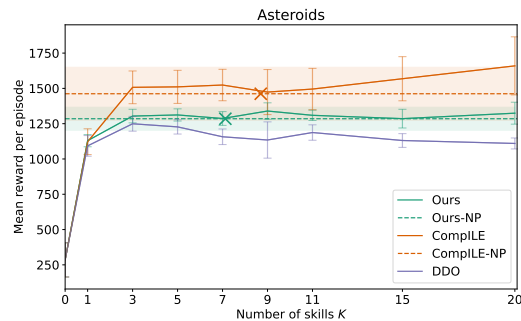


Figure 4. Additional Atari environment, same setup as in Figure 3.

specifying K as a hyperparameter can be avoided by placing a GEM prior over skills, both in our model and other skill discovery frameworks. We hope our work will motivate further research combining reinforcement learning and Bayesian nonparametrics, for example through small-variance asymptotics (Kulis & Jordan, 2011; Jiang et al., 2012; Broderick et al., 2013; Roychowdhury et al., 2013) for hard rather than soft clustering of skills, or with hierarchical models (Teh et al., 2006). Finally, we also hope that the technique we developed to increase K throughout training will find uses in other Bayesian nonparametric models, outside of skill discovery.

Acknowledgements

We thank the anonymous reviewers for their feedback, which helped improve our paper. We also thank Junfeng Wen for a useful suggestion on an ablation experiment. Our code was written in Python (vanRossum, 1995; Oliphant, 2007) and particularly relied on the following packages: Matplotlib (Hunter, 2007), TensorFlow (Abadi et al., 2015) (in particular for TensorBoard), gym (Brockman et al., 2016), Jupyter Notebook (Kluyver et al., 2016), PyTorch (Paszke et al., 2019), NumPy (Harris et al., 2020), and Stable-Baselines3 (Raffin et al., 2021).

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>. Software

- available from tensorflow.org.
- Achiam, J., Edwards, H., Amodei, D., and Abbeel, P. Variational option discovery algorithms. *arXiv preprint arXiv:1807.10299*, 2018.
- Ajay, A., Kumar, A., Agrawal, P., Levine, S., and Nachum, O. {OPAL}: Offline primitive discovery for accelerating offline reinforcement learning. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=V69LGwJ01IN>.
- Argall, B. D., Chernova, S., Veloso, M., and Browning, B. A survey of robot learning from demonstration. *Robotics and autonomous systems*, 57(5):469–483, 2009.
- Bacon, P.-L., Harb, J., and Precup, D. The option-critic architecture. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.
- Bellemare, M. G., Naddaf, Y., Veness, J., and Bowling, M. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- Broderick, T., Kulis, B., and Jordan, M. Mad-bayes: Map-based asymptotic derivations from bayes. In *International Conference on Machine Learning*, pp. 226–234. PMLR, 2013.
- Burda, Y., Grosse, R., and Salakhutdinov, R. Importance weighted autoencoders. *arXiv preprint arXiv:1509.00519*, 2015.
- Caron, F., Davy, M., and Doucet, A. Generalized polya urn for time-varying dirichlet process mixtures. In *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence*, pp. 33–40, 2007.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22, 1977.
- Dietterich, T. G. et al. The maxq method for hierarchical reinforcement learning. In *ICML*, volume 98, pp. 118–126. Citeseer, 1998.
- Dunson, D. B. and Xing, C. Nonparametric bayes modeling of multivariate categorical data. *Journal of the American Statistical Association*, 104(487):1042–1051, 2009.
- Esmaili, N., Sammut, C., and Shirazi, G. Behavioural cloning in control of a dynamic system. In *1995 IEEE International Conference on Systems, Man and Cybernetics. Intelligent Systems for the 21st Century*, volume 3, pp. 2904–2909. IEEE, 1995.
- Eysenbach, B., Gupta, A., Ibarz, J., and Levine, S. Diversity is all you need: Learning skills without a reward function. *arXiv preprint arXiv:1802.06070*, 2018.
- Ferguson, T. S. A bayesian analysis of some nonparametric problems. *The annals of statistics*, pp. 209–230, 1973.
- Florensa, C., Duan, Y., and Abbeel, P. Stochastic neural networks for hierarchical reinforcement learning. *arXiv preprint arXiv:1704.03012*, 2017.
- Fox, R., Krishnan, S., Stoica, I., and Goldberg, K. Multi-level discovery of deep options. *arXiv preprint arXiv:1703.08294*, 2017.
- Gershman, S. and Goodman, N. Amortized inference in probabilistic reasoning. In *Proceedings of the annual meeting of the cognitive science society*, volume 36, 2014.
- Glynn, P. W. Likelihood ratio gradient estimation for stochastic systems. *Communications of the ACM*, 33(10):75–84, 1990.
- Gregor, K., Rezende, D. J., and Wierstra, D. Variational intrinsic control. *arXiv preprint arXiv:1611.07507*, 2016.
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., del Río, J. F., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., and Oliphant, T. E. Array programming with NumPy. *Nature*, 585(7825):357–362, September 2020. doi: 10.1038/s41586-020-2649-2. URL <https://doi.org/10.1038/s41586-020-2649-2>.
- Ho, J. and Ermon, S. Generative adversarial imitation learning. *Advances in neural information processing systems*, 29:4565–4573, 2016.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9:1735–1780, 1997.
- Horgan, D., Quan, J., Budden, D., Barth-Maron, G., Hessel, M., Van Hasselt, H., and Silver, D. Distributed prioritized experience replay. *arXiv preprint arXiv:1803.00933*, 2018.
- Hunter, J. D. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007. doi: 10.1109/MCSE.2007.55.

- Hussein, A., Gaber, M. M., Elyan, E., and Jayne, C. Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)*, 50(2):1–35, 2017.
- Ishwaran, H. and James, L. F. Gibbs sampling methods for stick-breaking priors. *Journal of the American Statistical Association*, 96(453):161–173, 2001.
- Jang, E., Gu, S., and Poole, B. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- Jiang, K., Kulis, B., and Jordan, M. Small-variance asymptotics for exponential family dirichlet process mixture models. *Advances in Neural Information Processing Systems*, 25:3158–3166, 2012.
- Khetarpal, K., Klissarov, M., Chevalier-Boisvert, M., Bacon, P.-L., and Precup, D. Options of interest: Temporal abstraction with interest functions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 4444–4451, 2020.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Kipf, T., Li, Y., Dai, H., Zambaldi, V., Sanchez-Gonzalez, A., Grefenstette, E., Kohli, P., and Battaglia, P. Compile: Compositional imitation learning and execution. In *International Conference on Machine Learning*, pp. 3418–3428. PMLR, 2019.
- Kluyver, T., Ragan-Kelley, B., Pérez, F., Granger, B., Bussonnier, M., Frederic, J., Kelley, K., Hamrick, J., Grout, J., Corlay, S., Ivanov, P., Avila, D., Abdalla, S., and Willing, C. Jupyter notebooks – a publishing format for reproducible computational workflows. In Loizides, F. and Schmidt, B. (eds.), *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, pp. 87 – 90. IOS Press, 2016.
- Koller, D. and Friedman, N. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- Kool, W., van Hoof, H., and Welling, M. Estimating gradients for discrete random variables by sampling without replacement. *arXiv preprint arXiv:2002.06043*, 2020.
- Krishnan, S., Garg, A., Patil, S., Lea, C., Hager, G., Abbeel, P., and Goldberg, K. Transition state clustering: Unsupervised surgical trajectory segmentation for robot learning. In *Robotics Research*, pp. 91–110. Springer, 2018.
- Krishnan, S., Garg, A., Liaw, R., Thananjeyan, B., Miller, L., Pokorny, F. T., and Goldberg, K. Swirl: A sequential windowed inverse reinforcement learning algorithm for robot tasks with delayed rewards. *The International Journal of Robotics Research*, 38(2-3):126–145, 2019.
- Kulis, B. and Jordan, M. I. Revisiting k-means: New algorithms via bayesian nonparametrics. *arXiv preprint arXiv:1111.0352*, 2011.
- Maddison, C. J., Mnih, A., and Teh, Y. W. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *nature*, 518(7540): 529–533, 2015.
- Murali, A., Garg, A., Krishnan, S., Pokorny, F. T., Abbeel, P., Darrell, T., and Goldberg, K. Tsc-dl: Unsupervised trajectory segmentation of multi-modal surgical demonstrations with deep learning. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4150–4157. IEEE, 2016.
- Nalisnick, E. and Smyth, P. Stick-breaking variational autoencoders. *arXiv preprint arXiv:1605.06197*, 2016.
- Neal, R. M. Markov chain sampling methods for dirichlet process mixture models. *Journal of computational and graphical statistics*, 9(2):249–265, 2000.
- Niekum, S., Chitta, S., Barto, A. G., Marthi, B., and Osentoski, S. Incremental semantically grounded learning from demonstration. In *Robotics: Science and Systems*, volume 9, pp. 10–15607. Berlin, Germany, 2013.
- Oliphant, T. E. Python for scientific computing. *Computing in science & engineering*, 9(3):10–20, 2007.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32:8026–8037, 2019.
- Paulus, M. B., Maddison, C. J., and Krause, A. Rao-blackwellizing the straight-through gumbel-softmax gradient estimator. *arXiv preprint arXiv:2010.04838*, 2020.
- Potapczynski, A., Loaiza-Ganem, G., and Cunningham, J. P. Invertible gaussian reparameterization: Revisiting the gumbel-softmax. *Advances in Neural Information Processing Systems*, 33, 2020.
- Raffin, A., Hill, A., Gleave, A., Kanervisto, A., Ernestus, M., and Dormann, N. Stable-baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, 22(268):1–8, 2021. URL <http://jmlr.org/papers/v22/20-1364.html>.

- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. In *International conference on machine learning*, pp. 1278–1286. PMLR, 2014.
- Roeder, G., Wu, Y., and Duvenaud, D. K. Sticking the landing: Simple, lower-variance gradient estimators for variational inference. *Advances in Neural Information Processing Systems*, 30:6925–6934, 2017.
- Ross, S., Gordon, G., and Bagnell, D. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 627–635. JMLR Workshop and Conference Proceedings, 2011.
- Roychowdhury, A., Jiang, K., and Kulis, B. Small-variance asymptotics for hidden markov models. In *Advances in Neural Information Processing Systems*, pp. 2103–2111, 2013.
- Schaal, S. et al. Learning from demonstration. *Advances in neural information processing systems*, pp. 1040–1046, 1997.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Sethuraman, J. A constructive definition of dirichlet priors. *Statistica sinica*, pp. 639–650, 1994.
- Shankar, T. and Gupta, A. Learning robot skills with temporal variational inference. In *International Conference on Machine Learning*, pp. 8624–8633. PMLR, 2020.
- Sharma, A., Sharma, M., Rhinehart, N., and Kitani, K. M. Directed-info gail: Learning hierarchical policies from unsegmented demonstrations using directed information. *arXiv preprint arXiv:1810.01266*, 2018.
- Sharma, A., Gu, S., Levine, S., Kumar, V., and Hausman, K. Dynamics-aware unsupervised discovery of skills. *arXiv preprint arXiv:1907.01657*, 2019.
- Shiarlis, K., Wulfmeier, M., Salter, S., Whiteson, S., and Posner, I. TACO: Learning task decomposition via temporal alignment for control. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4654–4663. PMLR, 10–15 Jul 2018. URL <https://proceedings.mlr.press/v80/shiarlis18a.html>.
- Stirn, A., Jebara, T., and Knowles, D. A new distribution on the simplex with auto-encoding applications. *Advances in Neural Information Processing Systems*, 32:13670–13680, 2019.
- Such, F. P., Madhavan, V., Liu, R., Wang, R., Castro, P. S., Li, Y., Zhi, J., Schubert, L., Bellemare, M. G., Clune, J., et al. An atari model zoo for analyzing, visualizing, and comparing deep reinforcement learning agents. *Proceedings of IJCAI 2019*, 2019. URL <https://github.com/uber-research/atari-model-zoo>.
- 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., Jordan, M. I., Beal, M. J., and Blei, D. M. Hierarchical dirichlet processes. *Journal of the american statistical association*, 101(476):1566–1581, 2006.
- vanRossum, G. Python reference manual. *Department of Computer Science [CS]*, (R 9525), 1995.
- 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*, pp. 3540–3549. PMLR, 2017.
- Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256, 1992.

A. Relaxations

We choose the following relaxations:

$$\delta_{b_0^{(i)}=1} \approx b_0^{(i)} \quad (23)$$

$$\delta_{h_t^{(i)}=h_{t-1}^{(i)}} \approx 1 - \|h_t^{(i)} - h_{t-1}^{(i)}\|_1/2 \quad (24)$$

$$\eta(h_t^{(i)}) \approx \eta^> h_t^{(i)} \quad (25)$$

$$p_\theta(b_t^{(i)}, h_t^{(i)} | h_{t-1}^{(i)}, s_t^{(i)}) \approx b_t^{(i)} \psi_{h_{t-1}^{(i)}}(s_t^{(i)}) \eta^> h_t^{(i)} + (1 - b_t^{(i)}) (1 - \psi_{h_{t-1}^{(i)}}(s_t^{(i)})) \cdot (1 - \|h_t^{(i)} - h_{t-1}^{(i)}\|_1/2) \quad (26)$$

Note that these relaxations (1) match the objective they are relaxing when b 's are binary and h 's are one-hot; and (2) that they are not unique: there are many sensible choices that could be used.

B. Learning in Augmented Environments

In the augmented environment, taking an action corresponding to one of the skills means that the agent may actually interact with the environment for several timesteps. Naïvely, one might think that simply using the usual Bellman backup update considering these enhanced actions as actions would be sufficient. However, this neglects the fact that the reward credited to the enhanced action arises from multiple environment interactions, and that the discounting of the Q -value of the next state depends on the length of the enhanced action. Recall that the usual Bellman equation is:

$$Q(s, a) = \mathbb{E}_\pi \left[\sum_{t=1}^{\infty} \gamma^{t-1} R_t | S_0 = s, A_0 = a \right] \quad (27)$$

$$= \mathbb{E}_\pi [R_1 + \gamma Q(S_1, A_1) | S_0 = s, A_0 = a] \quad (28)$$

where the first term in the last equality (R_1) corresponds to the reward from a single timestep, and the second term ($\gamma Q(S_1, A_1)$) to the discounted return of all the following states.

For an augmented environment, let action a result in τ interactions with the (unaugmented) environment. Note that τ is a random variable, since skill terminations are stochastic, though of course $\tau = 1$ deterministically if a is a primitive action. We then have:

$$Q(s, a) = \mathbb{E}_{\pi, \tau} \left[\sum_{t=1}^{\tau} \gamma^{t-1} R_t + \gamma^\tau Q(S_\tau, A_\tau) | S_0 = s, A_0 = a \right] \quad (29)$$

As before, t indexes interactions with the unaugmented environment during execution of a single enhanced action. The first term ($\sum_{t=1}^{\tau} \gamma^{t-1} R_t$) corresponds to the reward directly credited to the (possibly enhanced) action, and the second term ($\gamma^\tau Q(S_\tau, A_\tau)$) is the discounted expected return from future (possibly enhanced) actions. As a sanity check, note that in the case where τ is always 1, both equations match.

This modified Bellman equation implies that we must keep a slightly modified replay buffer. The typical buffer $\{(s_t, a_t, s_{t+1}, r_t)\}_t$, storing every interaction with the unaugmented environment, is insufficient. We would be unable to determine when enhanced actions were taken or when they terminated, and therefore unable to credit them accordingly. Further, such a buffer is wasteful, assuming that enhanced actions are common. Instead we should keep a buffer $\{(s_t, a_t, s_{t+\tau_t}, \sum_{t'=t}^{t+\tau_t} \gamma^{t'-t} r_{t'}, \tau_t)\}$. Here a_t denotes the action taken in the augmented action space at time t (possibly a primitive action, possibly an enhanced action). τ_t is the (random) duration of that action. We must keep τ_t to be able to discount Q on the right hand side of the Bellman equation appropriately. We need only put entries into the buffer at timesteps at which we take an action in the augmented environment. We do not need to record every interaction with the unaugmented environment.

C. Experimental Details

Proof-of-concept environment For this experiment, we generate 1000 expert trajectories using a manually designed policy. We then train our model for 500 epochs. The options' sub-policies and termination functions consist of MLPs with two hidden layers of 16 units separated by a ReLU activation and followed by a Softmax activation. The parameters of all

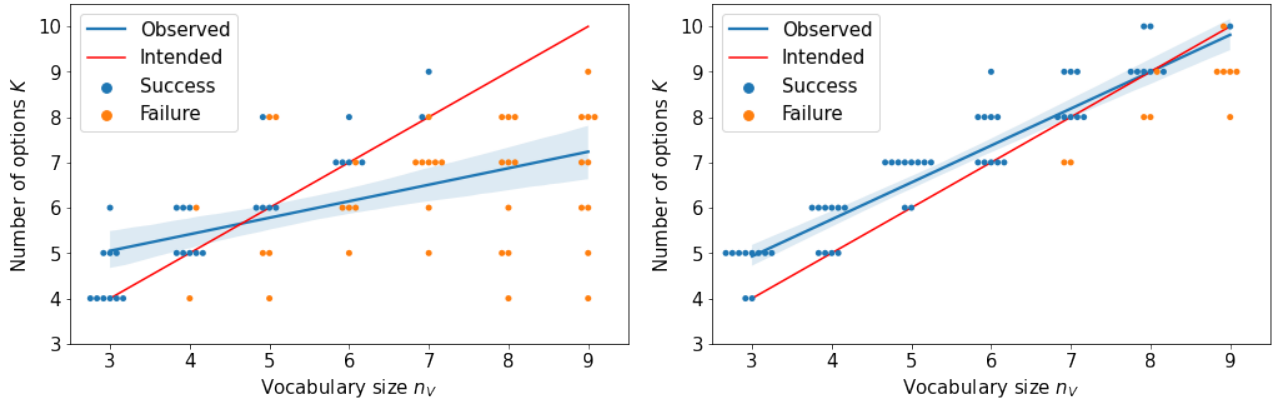


Figure 5. Results on our proof-of-concept environment, without entropy regularization (left panel), and with fixed entropy regularization (right panel).

layers except the last one are shared across options. For the encoder, we use an LSTM layer with 32 hidden units, and MLPs with two hidden layers of 32 units and ReLU activations for both heads, with weights shared except for the last layer. We use a learning rate of 0.005 with the Adam optimizer (Kingma & Ba, 2014) and a batch size of 128. The GS temperature parameter is initialized at 1 and annealed by a factor of 0.995 each epoch. λ_{ent} is initialized at 5 and also annealed by a factor of 0.995 each epoch. We check the option usage every 10 epochs ($n_K = 10^4$) and use $\delta = 0.5$ for our rule of increasing K .

Atari environments For these experiments, we use 1000 expert trajectories of length 300 and train for 1500 epochs. We use the RAM state of the Atari environments, where each state is a 128 bytes array, that we unpack into a 1024 bit array. The options’ sub-policies and termination functions consist of a single linear layer with a Softmax activation per option. This choice is motivated by the fact that we do not want a single sub-policy to be able to fully reconstruct the expert trajectories. We use the same encoder as for the proof-of-concept environment, and the same optimizer with a learning rate of 0.001. The GS temperature is now annealed by a factor of 0.999. We keep the same values for n_K and δ . We use similar architectures for sub-policies and termination in both CompILE and DDO, as well as for the encoder in CompILE. We fix the number of segments in CompILE to 7 for all runs.

To learn in the augmented environment, we use the PPO agent implemented by Raffin et al. (2021), with the default ‘CnnPolicy’ that takes as input the image state of the environment with the same preprocessing as done by Mnih et al. (2015). We use the implementation default parameters except for the ‘n_steps’ variable (the number of environment steps used per update) that we set to 512. We also modify the replay buffer used during training to take into account the specific aspects of learning in an augmented environment mentioned in Appendix B. This agent is trained for $3 \cdot 10^6$ (augmented) steps and the mean reward across the last $5 \cdot 10^5$ steps is used in Fig. 3. For CompILE, we automatically terminate each enhanced action after 15 time steps for all augmented environments, as doing so was preferable to following the Poisson-sampled termination.

D. Additional Experiments

D.1. Proof-of-Concept Environment

We show in Figure 5 results analogous to those of Figure 2, except we do not use the entropy regularizer from Section 4.4 (left panel), or we simply do not anneal it (right panel). We can see that, as mentioned in the main manuscript, not using the regularizer significantly degrades performance, although not using annealing (and keeping the regularizing coefficient fixed throughout training) does not have much of an impact.

D.2. Atari Environments

As mentioned in the main manuscript, we ablate some of our choices. Figure 6 shows our ablation results of using the entropy regularizer, green curves show results with the regularizer, and blue ones without. Across environments, performance is

comparable, so that the entropy regularizer is not truly needed here. We highlight this result is opposite to what we found in our proof-of-concept environment, where the regularizer was fundamental to getting our method to work. Since adding the regularizer does not hurt performance in Atari environments and greatly helps in our proof-of-concept environment, we nonetheless recommend to use it as a default.

Tables 1, 2, and 3 show the results of ablations where the length of expert trajectories is changed for a subset of the Atari environments that we considered. We highlight that we did not cherry pick these environments, and the fact that we do not show analogous results for the missing environments was merely a matter of computational costs. While results do change significantly by varying trajectory length, both for our methods and the baselines, we can see that: (1) our fixed- K method consistently outperforms or remains competitive with both CompILE and DDO, the only exception being Montezuma’s revenge with trajectories of length 50, and (2) our nonparametric method outperforms or remains competitive with nonparametric CompILE across all settings. We also highlight that DDO ran out of memory when using trajectories of length 1000. We can thus see that our empirical superiority shown in the main manuscript was not due to a lucky choice of expert trajectory length.

Table 1. Reward on Montezuma’s revenge, varying the length of expert trajectories. Columns requiring K , i.e. “Ours”, “CompILE”, and “DDO”, use $K = 7$.

Length	Ours	CompILE	DDO	Ours-NP	CompILE-NP
50	53.8 ± 7.6	113.1 ± 65.6	41.8 ± 20.9	274.0 ± 51.6	149.9 ± 73.7
150	350.9 ± 44.8	0.3 ± 0.3	24.3 ± 19.2	344.9 ± 44.6	102.7 ± 91.7
300	317.4 ± 48.6	0.0 ± 0.0	0.3 ± 0.2	261.0 ± 38.8	92.6 ± 80.0
500	350.9 ± 42.5	0.0 ± 0.0	27.0 ± 0.0	391.2 ± 31.8	55.1 ± 48.4
1000	328.3 ± 86.2	0.0 ± 0.0	NA	610.9 ± 74.4	66.6 ± 45.0

Table 2. Reward on Breakout, varying the length of expert trajectories. Columns requiring K use $K = 7$ for “Ours” and “CompILE” and $K = 11$ for “DDO”.

Length	Ours	CompILE	DDO	Ours-NP	CompILE-NP
50	32.8 ± 1.0	16.0 ± 2.9	27.4 ± 1.3	31.6 ± 1.5	25.9 ± 2.6
150	27.1 ± 1.3	20.0 ± 1.2	25.3 ± 1.9	22.9 ± 1.1	21.5 ± 1.1
300	36.6 ± 3.0	18.6 ± 2.1	26.5 ± 1.7	31.6 ± 2.9	27.4 ± 2.9
500	28.0 ± 2.0	18.4 ± 1.6	20.6 ± 0.0	22.6 ± 1.1	24.0 ± 3.4
1000	23.0 ± 1.7	16.4 ± 3.2	NA	19.8 ± 1.1	17.9 ± 0.9

Table 3. Reward on Space Invaders, varying the length of expert trajectories. Columns requiring K , i.e. “Ours”, “CompILE”, and “DDO”, use $K = 7$.

Length	Ours	CompILE	DDO	Ours-NP	CompILE-NP
50	488.1 ± 10.5	367.7 ± 18.1	414.1 ± 15.6	418.2 ± 26.6	440.2 ± 17.8
150	562.0 ± 21.6	401.2 ± 14.8	475.0 ± 17.3	492.0 ± 16.1	447.3 ± 20.3
300	583.0 ± 16.2	438.8 ± 45.6	479.3 ± 19.9	531.0 ± 17.7	469.9 ± 15.9
500	562.9 ± 18.9	432.6 ± 29.2	468.8 ± 27.1	478.4 ± 12.9	469.2 ± 17.2
1000	552.5 ± 19.4	420.8 ± 24.4	NA	490.8 ± 16.3	502.3 ± 25.5

Bayesian Nonparametrics for Offline Skill Discovery

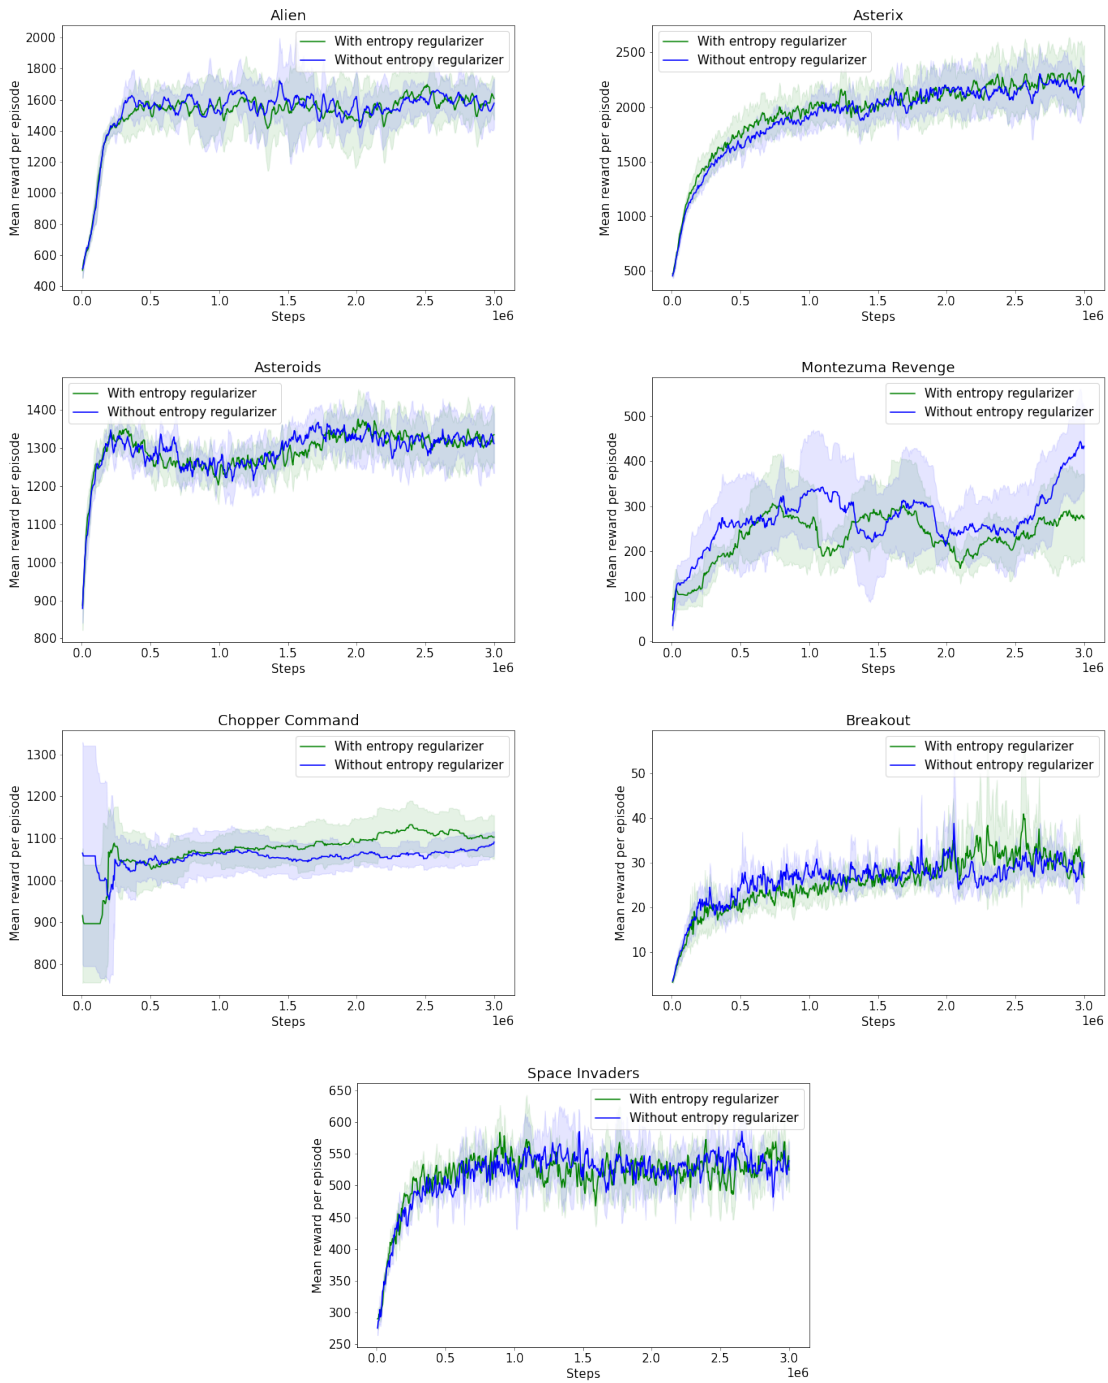


Figure 6. Evolution of the mean reward per episode during training with and without the entropy regularizer. Results are averaged across 5 runs.