

---

# Variational Empowerment as Representation Learning for Goal-Based Reinforcement Learning

---

Jongwook Choi<sup>†1</sup> Archit Sharma<sup>†2</sup> Honglak Lee<sup>13</sup> Sergey Levine<sup>45</sup> Shixiang Shane Gu<sup>4</sup>

## Abstract

Learning to reach goal states and learning diverse skills through mutual information (MI) maximization have been proposed as principled frameworks for self-supervised reinforcement learning, allowing agents to acquire broadly applicable multi-task policies with minimal reward engineering. Starting from a simple observation that the standard goal-conditioned RL (GCRL) is encapsulated by the optimization objective of variational empowerment, we discuss how GCRL and MI-based RL can be generalized into a single family of methods, which we name *variational GCRL* (*VGCR*L), interpreting variational MI maximization, or variational empowerment, as representation learning methods that acquire functionally-aware state representations for goal reaching. This novel perspective allows us to: (1) derive simple but unexplored variants of GCRL to study how adding small representation capacity can already expand its capabilities; (2) investigate how discriminator function capacity and smoothness determine the quality of discovered skills, or latent goals, through modifying latent dimensionality and applying spectral normalization; (3) adapt techniques such as hindsight experience replay (HER) from GCRL to MI-based RL; and lastly, (4) propose a novel evaluation metric, named latent goal reaching (LGR), for comparing empowerment algorithms with different choices of latent dimensionality and discriminator parameterization. Through principled mathematical derivations and careful experimental studies, our work lays a novel foundation from which to evaluate, analyze, and develop representation learning techniques in goal-based RL.

## 1 Introduction

Reinforcement learning (RL) provides a general framework for discovering optimal behaviors for sequential decision-making. Combined with powerful function approximators like neural networks, RL can be used to learn to play computer games from raw pixels (Mnih et al., 2013) and acquire complex sensorimotor skills with real-world robots (Gu et al., 2017a; Kalashnikov et al., 2018; Harnoja et al., 2018). Neural networks show best performance, generalization, and reusability when they are trained on large and diverse datasets (Krizhevsky et al., 2012; Devlin et al., 2018). However, a critical limitation in RL is that human experts often need to spend considerable efforts designing and fine-tuning reward functions per task, making it hard to scale and define a huge set of tasks in advance. If we have agents that can interact with the world without rewards, build up a body of knowledge autonomously, and utilize this knowledge to accomplish new tasks efficiently, then we can greatly scale up task and skill learning to achieve similar level of generalization and performance for RL as what neural networks have enabled for other domains.

Several works have tried to find a single generalizable task-agnostic reward function which can potentially be used across several environments. The use of such intrinsic reward functions has been motivated as exploration heuristics such as curiosity and novelty (Schmidhuber, 1991; Oudeyer & Kaplan, 2009; Bellemare et al., 2016; Pathak et al., 2017), as optimizing mutual information (MI) (Gregor et al., 2017; Eysenbach et al., 2019; Sharma et al., 2020b;a) or as empowerment (Klyubin et al., 2005; Jung et al., 2011; Mohamed & Rezende, 2015). Classically, goal-conditioned RL (GCRL) has shown success in learning diverse and useful skills in concurrence to MI-based methods. GCRL optimizes a stationary and interpretable reward for goal-reaching, but when the goal space is high-dimensional, how does the agent know which part of the space is relevant and which part can be ignored? In such cases, prior GCRL works frequently rely on manual definition (Andrychowicz et al., 2017) or off-the-shelf representation learning (Nachum et al., 2018; Nair et al., 2018; Wu et al., 2018) optimized prior to or separately from reinforcement learning. Meanwhile, MI or empowerment-based RL offers a clear objective for representation learning *through* reinforcement learning, but the

---

<sup>†</sup>Work done while an intern at Google.

<sup>1</sup>University of Michigan <sup>2</sup>Stanford University <sup>3</sup>LG AI Research

<sup>4</sup>Google Research <sup>5</sup>University of California, Berkeley. Correspondence to: Jongwook Choi <jwook@umich.edu>, Shixiang Shane Gu <shanegu@google.com>.

properties of the learned behaviors are often unclear due to lack of a proper evaluation metric. Prior works use qualitative inspections of learned behaviors, variational bound estimates, or downstream task performances of a skill-utilizing high-level policy (Eysenbach et al., 2019; Sharma et al., 2020b), but these heuristics are costly or indirect measures and make objective comparisons and analyses of various mathematically-similar MI-based algorithms difficult (Florensa et al., 2017; Eysenbach et al., 2019; Achiam et al., 2018; Warde-Farley et al., 2019; Hansen et al., 2020; Sharma et al., 2020b). To recover a more direct metric, an important question is: what do these MI-based objectives learn representations for?

In this work, we interpret MI and empowerment-based RL as a principled framework for representation learning in goal-conditioned RL. Starting from a simple observation that the objective of the standard GCRL can be seen as a special case of variational MI with a fixed hard-coded variational posterior, our analysis provides a unification of these ideas and explicitly reframes skill discovery via mutual information maximization (Gregor et al., 2017; Eysenbach et al., 2019) as a combination of representation learning and goal-conditioned reinforcement learning, where both the space of goals and the skills to reach those goals are learned jointly via a MI-based objective. While the connections between representation learning, mutual information estimation, and goal-conditioned RL have been explored in a number of previous works (Gregor et al., 2017; Warde-Farley et al., 2019; Gupta et al., 2018), our exact mathematical formulation and granular analyses enable new perspectives and synergies between GCRL and MI-based RL:

1. **[MI to GCRL]** We propose simple but novel variants of GCRL – adaptive-variance and linear-mapping GCRL – to study how adding small representation capacity can already expand the capabilities of GCRL.
2. **[MI to GCRL]** We show that a proper representation regularization from generative modeling, such as spectral normalization (Miyato et al., 2018), can improve the quality of latent goals discovered (and the stability of MI-based algorithms).
3. **[GCRL to MI]** We adapt hindsight experience replay (HER) (Andrychowicz et al., 2017) from GCRL to more general MI-based objectives and show posterior HER (P-HER) consistently provides substantial performance gains in MI-based RL algorithms.
4. **[GCRL to MI]** We propose the latent goal reaching (LGR) metric as an intuitive, task-oriented, and discriminator-agnostic metric for objectively evaluating empowerment algorithms.

## 2 Related Work

Reward engineering has been a bottleneck to broad application of RL. Some of the prior attempts to alleviate this problem have sought to introduce human supervision in alternative, easier forms, such as demonstrations (Ng et al., 2000; Abbeel & Ng, 2004; Ziebart et al., 2008; Ho & Ermon, 2016; Fu et al., 2017; Ghasemipour et al., 2019) or preferences (Hadfield-Menell et al., 2017; Christiano et al., 2017). However, since these methods still rely on non-negligible amounts of human interventions, they cannot automatically scale to solving thousands of new environments and tasks.

**Empowerment and reward-free RL.** Task-agnostic reward functions have been proposed to encourage exploration in environments using notions of curiosity or novelty (Schmidhuber, 1991; Oudeyer & Kaplan, 2009; Schmidhuber, 2010; Bellemare et al., 2016; Pathak et al., 2017; Colas et al., 2018). In a similar vein, some methods maximize the state-visitation entropy (Hazan et al., 2018; Pong et al., 2019; Lee et al., 2019; Ghasemipour et al., 2019). These approaches can enable solutions to otherwise hard exploration sparse-reward problems. Some of the recent work has emphasized on empowerment or option/skill discovery through optimization of mutual information based intrinsic reward functions. Classically, empowerment measures the ability of an agent to control the environment (Salge et al., 2014; Klyubin et al., 2005; Jung et al., 2011), which was scaled up by (Mohamed & Rezende, 2015; Karl et al., 2017). The concept of mutual information, which is also at the heart of empowerment based methods, has been further used to motivate several objectives for skill discovery (Florensa et al., 2017; Eysenbach et al., 2019; Achiam et al., 2018; Warde-Farley et al., 2019; Hansen et al., 2020; Sharma et al., 2020b; Campos et al., 2020). Recent works have shown that skills learned through mutual information can be meaningfully combined to solve downstream tasks (Eysenbach et al., 2019; Sharma et al., 2020b), even on real robots (Sharma et al., 2020a).

**Goal-conditioned RL.** Goal-conditioned RL (Kaelbling, 1993; Pong et al., 2018; Andrychowicz et al., 2017; Schaul et al., 2015) provides a framework for enabling agents to reach user-specified goal states. The behaviors learned via GCRL can be interpretable and easy to analyze in terms of the goal-reaching function. However, GCRL assumes that a goal-reaching function has been specified in addition to goal states, which precludes broader application for the same reasons as those for reward engineering. To overcome such limitations, some prior works have used variational inference (Rudner et al., 2021) or mutual information in the GCRL framework (Pong et al., 2019; Warde-Farley et al., 2019) to learn. However, these works use the MI optimization as an unsupervised scheme to generate and achieve goals. On the other hand, our work studies skill-discovery/empowerment methods and provides an explicit

reinterpretation within the GCRL framework, combining the representation learning perspective with the goal-reaching behavior of GCRL.

### 3 Background

In this section, we briefly review mutual information (MI)-based objectives for skill discovery, focusing on variational approaches introduced in (Mohamed & Rezende, 2015; Gregor et al., 2017; Eysenbach et al., 2019; Sharma et al., 2020b), and goal-conditioned RL (GCRL).

We denote a Markov decision process (MDP)  $\mathcal{M} = (\mathcal{S}, \mathcal{A}, p, r)$ , where  $\mathcal{S}$  denotes the state space,  $\mathcal{A}$  denotes the action space,  $p : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \rightarrow [0, \infty)$  denotes the underlying (possibly stochastic) transition dynamics of the environment with the initial state distribution  $p_0 : \mathcal{S} \rightarrow [0, \infty)$ , and a reward function  $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ . The goal of the RL optimization problem is to learn a policy  $\pi(a | s)$  which maximizes the return  $\mathbb{E}_{p, \pi} [\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t)] = \mathbb{E}_{s \sim \rho^\pi, a \sim \pi} [r(s, a)]$ , for a discount factor  $\gamma \in [0, 1)$  where  $\rho^\pi$  is an unnormalized  $\gamma$ -discounted state visitation density. Importantly, once we write an objective in the form of  $\mathbb{E}_{\rho^\pi, \pi} [r(s, a)]$ , we can apply the policy gradient theorem (Sutton et al., 2000) to derive a practical RL solver, as done in (Kakade, 2002; Silver et al., 2014; Schulman et al., 2015; Gu et al., 2017b; Ciosek & Whiteson, 2018), or learn it with Q-learning (Watkins & Dayan, 1992). For simplicity of our notations, we omit the discount factor  $\gamma$  in the following sections and derivations.

#### 3.1 Mutual Information Maximization and Empowerment

MI maximization in RL such as empowerment generally means maximizing the mutual information between (some representations of) actions and (some representations of) future states following those actions (Klyubin et al., 2005; Mohamed & Rezende, 2015). The goal is to learn a set of actions that can influence future states to be diverse, but also be predictable if we know what action is taken. In this work we focus on learning abstract representation  $z \in \mathcal{Z}$ , which is an additional input to the policy  $\pi(a | s, z)$  and defines a set of empowered actions. The latent code  $z$  (either discrete or continuous) can be interpreted as a macro-action, *skill* or *goal* (Eysenbach et al., 2019; Sharma et al., 2020b).

We discuss two variants of MI objectives in RL: *state-predictive MI* (Sharma et al., 2020b), which maximizes  $\mathcal{I}(s'; z | s)$ , and *state-marginal MI* (Eysenbach et al., 2019), which maximizes  $\mathcal{I}(s; z)$ . Due to page limit, we discuss these variants more in detail in Appendix A. In this work, we focus on state-marginal MI, whose optimization objective is:

$$\begin{aligned} \mathcal{I}(s; z) &= \mathbb{E}_{z \sim p(z), s \sim \rho^\pi(s|z)} [\log p(z | s) - \log p(z)] \\ &\geq \mathbb{E}_{z \sim p(z), s \sim \rho^\pi(s|z)} [\log q_\lambda(z | s) - \log p(z)] \end{aligned} \quad (1)$$

where  $q_\lambda(z | s)$  is a variational approximation to the in-

tractable posterior  $p(z | s)$ , often called a (skill) discriminator (Eysenbach et al., 2019).

Given a parameterized policy  $\pi_\theta(a | s, z)$ , Eq. 1 gives a joint maximization objective (a variational lower bound) with respect to  $\pi_\theta$  and  $q_\lambda$ :

$$\mathcal{F}(\theta, \lambda) = \mathbb{E}_{z, s \sim \pi_\theta} [\log q_\lambda(z | s) - \log p(z)]. \quad (2)$$

A simple iterative RL procedure can be derived to optimize this lower bound, assuming a parameterized policy  $\pi_\theta(a | s, z)$ , where at iteration  $i$ ,

$$\lambda^{(i)} \leftarrow \operatorname{argmax}_\lambda \mathbb{E}_{z, s \sim \pi^{(i-1)}} [\log q_\lambda(z | s) - \log p(z)] \quad (3)$$

$$\theta^{(i)} \leftarrow \operatorname{argmax}_\theta \mathbb{E}_{z, s \sim \pi_\theta} [\log q_{\lambda^{(i)}}(z | s) - \log p(z)]. \quad (4)$$

Eq. 3 is a simple supervised regression (e.g., maximum likelihood) on on-policy samples. Eq. 4 has the same form of standard RL, and therefore can be optimized using any RL algorithm (Gregor et al., 2017; Eysenbach et al., 2019).

#### 3.2 Goal-Conditioned RL

Goal-conditioned RL (Kaelbling, 1993; Schaul et al., 2015) (GCRL) is a standard, stationary-reward problem where we aim to find a policy  $\pi(a | s, g)$  conditional on a goal  $g \in \mathcal{G}$  by maximizing

$$F(\pi) = \mathbb{E}_{g \sim p(g), s \sim \pi_\theta} [-d(s, g)], \quad (5)$$

where  $p(g, s) = p(g)\rho^\pi(s|g)$ ,  $p(g)$  defines the task distribution over goals, and  $d(s, g)$  is a distance metric between state  $s$  and goal  $g$ , such as an Euclidean distance. The main challenges for goal-conditioned RL lie in defining the goal space and the goal-reaching reward function  $-d(s, g)$ , which often requires task-specific knowledge or careful choices of goal space (Plappert et al., 2018).

In off-policy learning, hindsight experience replay (HER) (Kaelbling, 1993; Andrychowicz et al., 2017) has shown to improve learning of goal-conditioned policy significantly. The key insight is that for a given exploration episode  $\{g, s_{0:T}\}$ , one can relabel the goal with an *actually* achieved goal  $\mathbb{S}(s_{0:T})$ , derived by a strategy function  $\mathbb{S}(\cdot)$ . A typical choice is to relabel the goal as  $\tilde{g} = \mathbb{S}(s_{0:T}) = s_T$ , which can be seen as self-supervised curriculum learning (Andrychowicz et al., 2017; Lynch et al., 2019).

### 4 Expressivity Tradeoffs in Variational Empowerment

Interestingly, the simple objective in Eq. 2, which we term *Variational Goal-Conditioned RL* (VGCRRL), encapsulates most of the prior MI-based algorithms (Eysenbach et al., 2019; Warde-Farley et al., 2019; Hansen et al., 2020) with the only differences being goal space  $\mathcal{Z}$ , prior  $p(z)$ , and discriminator  $q_\lambda(z | s)$ , as detailed in Table 1. For example, when  $z$  is a discrete variable, this reduces to DIAYN (Eysenbach et al., 2019) or VALOR (Achiam et al., 2018).

Method	Goal space	$q_\lambda(z s)$	Learnable $\lambda$	Learning $\pi_z$
GCRL (Kaelbling, 1993)	Continuous ( $\mathbb{R}^d$ )	$\mathcal{N}(s, \sigma^2 I)$	-	(HER)
aGCRL (ours)	Continuous ( $\mathbb{R}^d$ )	$\mathcal{N}(s, \Sigma)$	$\Sigma$	HER
linGCRL (ours)	Continuous ( $\mathbb{R}^d$ )	$\mathcal{N}(As, \sigma^2 I)$	$A$	P-HER
InfoGAIL* (Li et al., 2017)	Discrete	Categorical	$q_\lambda$	-
DIAYN (Eysenbach et al., 2019)	Discrete	Categorical	$q_\lambda$	-
DIAYN (continuous)	Continuous ( $\mathbb{R}^d$ )	$\mathcal{N}(\mu(s), \Sigma(s))$	$\mu(\cdot)$	-
DISCERN (Warde-Farley et al., 2019)	$= \mathcal{S}$ (e.g. image)	Non-parametric	Embedding( $\cdot$ )	HER
VISR (Hansen et al., 2020)	Continuous ( $\mathbb{R}^d$ )	$\text{vMF}(\mu(s), \kappa)$	$\mu(\cdot)$	SF
VGCRl	Any	Any	Any	P-HER or SF

Table 1: A summary of algorithms, all are optimized with the single objective in Eq. 2. vMF stands for von Mises-Fisher distribution. DIAYN (Eysenbach et al., 2019), DISCERN (Warde-Farley et al., 2019), VISR (Hansen et al., 2020) can be seen as special cases. For InfoGAIL (Li et al., 2017)\*, we focus on the MI regularization objective  $L_I(\pi, Q)$  only. Since they are under the same objective, learning techniques for goal-conditioned policy  $\pi_z$  such as successor features (SF) (Barreto et al., 2017) and hindsight experience replay (HER) (Andrychowicz et al., 2017) can be adapted for more general settings within the VGCRl objective.

If  $z$  is continuous, a natural choice for  $q_\lambda$  is a Gaussian, i.e.  $\mathcal{N}(\mu(s), \Sigma(s))$ , where both  $\mu$  and  $\Sigma$  may be parameterized using any function approximators with a range of expressivities, from identity functions to deep neural networks. Throughout the rest of the paper, we show how various simple choices for  $q_\lambda$  lead to algorithms with different properties.

**Goal-Conditioned RL as a Coarse Variational Approximation.** A simple observation is that if we choose a fixed variational distribution, such as  $\mathcal{N}(s, \sigma^2 I)$  with  $\sigma$  as a fixed hyperparameter and the goal space identical to the observation space ( $\mathcal{Z} = \mathcal{S}$ ), the RL objective in Eq. 4 becomes (see Appendix B for mathematical details):

$$\mathcal{F}(\pi) = \mathbb{E}_{z, s \sim \pi_\theta} \left[ -\frac{1}{\sigma^2} \|z - s\|^2 \right] + \text{constant}. \quad (6)$$

It is straightforward to see that this recovers the objective of GCRL in Eq. 5 exactly (up to a constant), where the distance function uses a squared loss. This provides a novel interpretation for GCRL algorithms as a *variational empowerment algorithm with a hard-coded and fixed variational distribution*. Given that no  $q_\lambda$  parameters are adapted, this generally provides a very loose bound on MI; however, prior work on GCRL shows that this RL objective, unlike empowerment-based, learns useful goal-reaching skills stably (Kaelbling, 1993; Andrychowicz et al., 2017; Pong et al., 2018) thanks to a stationary reward function. This suggests that GCRL and prior variational empowerment methods represent two ends of a spectrum, corresponding to the expressivity of the variational distribution used to approximately maximize mutual information, and neither of the two is perfect, with their own pros and cons. Varying expressivity — through the choices of  $\mathcal{Z}$  and  $q_\lambda$  — and evaluating the qualities of learned goal spaces is a central theme of the next section.

## 5 Goal-Conditioned RL as Variational Empowerment

In this section, we discuss GCRL with representation learning, through the lens of variational empowerment: how the representation capacity leads to algorithms with different properties. We first derive two “lost relatives” of GCRL that only add minimal representation capacities but still result in interesting learning behaviors while keeping the stability of GCRL, and then discuss how we can study representation capacity in more general settings through varying smoothness constraints.

### 5.1 Adaptive Variances for Relevance Determination

Given the observation in Section 4, a straightforward modification to GCRL is to allow the variances to be learned, while keeping  $\mu(s) = s$ . If we assume a global learned covariance, i.e.,  $q_\lambda(z|s) = \mathcal{N}(s, \Sigma)$ ,  $\lambda = \{\Sigma\}$ , Eq. 2 gives us a novel variant of GCRL, which we call *adaptive GCRL* (aGCRL). The intuition behind this algorithm is the following: let us assume a simple diagonal covariance matrix; during learning, this algorithm will quickly shrink  $\sigma$  for the goal dimensions that the agent can reliably reach, and will expand variances for the dimensions that the agent has a hard time to; it therefore can identify and prioritize goal-reaching in feasible directions, discounting unfeasible ones, resembling properties of automatic relevance determination (ARD) (Wipf & Nagarajan, 2008).

**Experiment: Automatic Controllability Determination on Windy PointMass.** We design a simple *Windy Point-Mass* environment to study adaptive behaviors, which is simulated in Mujoco (Todorov et al., 2012). We assume a point mass in  $N$ -dimensional space (Figure 1a), where some dimensions have random force perturbations and therefore are difficult to control. Such perturbations are often studied in the risk-sensitive RL literature (Fox et al., 2015; Maddison et al., 2017); however, in our experiments, they

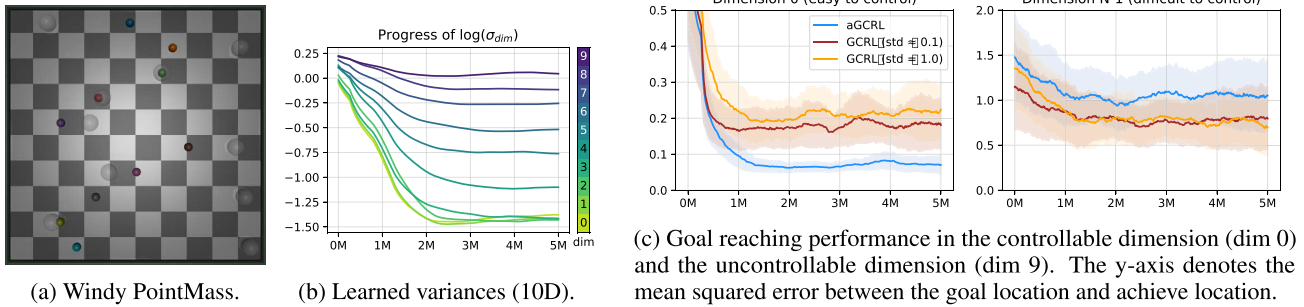


Figure 1: Adaptive-variance GCRL. The learned variance is clearly smaller for the easier (noiseless) dimension, which makes goal reaching in controllable dimensions more focused than in uncontrollable ones.

serve to create different levels of controllability. Our goal is to have GCRL automatically ignore dimensions that are not controllable and prioritize dimensions that are easy to control. More details can be found in Appendix D.

We evaluated goal-conditioned RL with an adaptive global diagonal variance term in Figure 1. Our results show that this simple modification to goal-based RL can accurately identify controllable dimensions in the state space. For example, in a 2-dimensional windy pointmass environment, aGCRL recovered a smaller variance for the first dimension,  $\sigma_x = 0.368$ , and a larger variance for the second dimension,  $\sigma_y = 1.648$ , which corresponds to having a reward function  $r(s, z) = \|x - g_x\|/0.368 + \|y - g_y\|/1.648$  where  $(x, y)$  is the position of the point mass and  $z = (g_x, g_y)$  is the goal location. A benefit of such adaptive variance is that we can prioritize goal reaching in controllable dimensions; in Figure 1c, we can observe that aGCRL can reach goals in the controllable dimension (e.g., dim 0) more quickly than the standard constant-variance GCRL baseline on the 10-dimensional Windy PointMass environment, showing the effectiveness of such *automatically learned reward functions that can ignore nuisance dimensions*.

## 5.2 Adaptive Mean with Varying Expressivity

The aGCRL variant adapts variances but fixes the mean  $\mu(s)$  to be  $s$ . By using more expressive parameterizations, such as neural networks (Eysenbach et al., 2019), the algorithm can theoretically optimize a tighter lower-bound to MI. However, as it gains more expressivity, interpretability and learning




Figure 2: linGCRL on 2D point mass. Left: the underlying physics space. Middle: An agent’s observation after a random linear projection. Right: A goal space recovered by  $\mu = Ao$ . The orange/green cross marks denote sampled goals  $z \sim p(z)$ , and blue dots are initial locations.

stability might be reduced. We study a linear case, i.e.  $q_\lambda(z|s) = \mathcal{N}(As, \Sigma)$  where  $\lambda = \{A\}$  along with identity and NN cases, and carefully evaluate this design choice.

**Experiment: Recovering Intrinsic Dimensions of Variations with Linear GCRL.** In this study, we design a simple 2D point mass with a random projection applied to the observation. We use an affine transformation  $W$  to generate the agent’s observation  $o = Ws$  from a physics simulator’s state  $s$ . For example, when a raw state in 2D point mass (Figure 2) includes the  $(x, y)$  location of the point mass, the agent will instead receive an entangled, obfuscated observation:  $o = (w_{11}x + w_{12}y, w_{21}x + w_{22}y, \dots)$  which does not align with the action space. We see whether a variant of VGCRL where  $q(z|o) = \mathcal{N}(Ao, \Sigma)$ , called *linGCRL*, can recover the inverse of an underlying projection  $A = W^{-1}$  when we use a rectangle-shaped 2D uniform prior  $p(z)$ . This resembles PCA discovering principal components in and underlying true dimensionalities.

Figure 2 shows an example where a unknown  $2 \times 2$  random projection  $W$  is applied. The arena of the 2D point mass environment and two example trajectories are visualized: the space of true state  $s$  (left), observation  $o = Ws$  (middle), and goal  $z = (A \cdot W)s$ . We can see two intrinsic, orthogonal dimensions are recovered by the learned matrix  $A$  such that the posteriors  $z$  from marginal states match the prior distribution, but up to rotation and reflection. This was also possible with more complex (e.g.,  $W \in \mathbb{R}^{10 \times 2}$ ) random projections. We show that linGCRL can *recover intrinsic dimensionalities of the state* from random projections.

## 5.3 Spectral Normalization

If the variational posterior (“discriminator”)  $q_\lambda(z|s)$  has high expressive power, for example when it is represented by a neural network, it can easily achieve maximum discriminability. However, this can lead to a very suboptimal solution for the policy. Because the marginal states as a result of latent goal and state pairs produced by the latent-conditioned policy are almost random and of poor quality in the early stage of training, the discriminator might easily *overfit* to such a near-random distribution of  $(s, z)$ ,




Figure 3: Visualization of the decision boundary of  $q(z|s)$  on 2D point mass environment with  $|G| = 20$  (See §5.3).

where  $s \sim \pi_\theta(\cdot|z)$ . This usually happens when in general  $q_\lambda(z|s)$  is much easier to fit (especially given a high capacity of neural networks) than the policy  $\pi_\theta$ . Figure 3(a) shows a motivating example of  $q_\lambda(z|s)$  learned on a toy 2D point mass environment: the landscape of  $q(z|s)$  is highly non-smooth, which can hinder learning latent-conditioned skills due to a ill-posed reward structure.

One way to alleviate this issue is to regularize the discriminator using Spectral Normalization (Miyato et al., 2018), which has been very effective in stabilizing Generative Adversarial Networks (GAN). Intuitively speaking, spectral normalization enforces the discriminator to be a *smooth* function by satisfying the Lipschitz continuity.

**Experiment: Spectral Normalization.** We study the behavior of variational empowerment algorithms on a toy 2D point mass environment, for a purpose of simple analysis. To simplify the analysis, we used  $\mu(s) = (x, y)$ , and a simple 2-layer MLP (with 128 hidden units). Figure 3 shows the decision boundary of the discriminator  $q_\lambda(z|s)$ , without and with Spectral Normalization (SN), where the dimension of  $z$  is 20 (or the number of skills). We can observe that, even though in both cases the empowerment objective  $\mathcal{F}$  has converged to near-optimum value with a reasonably good discriminability (Table 2: top-1 accuracy of  $q_\lambda(z|s)$  is  $> 90\%$ ), the decision boundary is much more smooth with Spectral Normalization and closer to interpretations of empowerment (Eysenbach et al., 2019; Mohamed & Rezende, 2015). Without Spectral Normalization, the discriminator had to learn a highly non-smooth decision boundary that is *over-fit* to near-random data generated the premature policy  $\pi_z$ , which would make joint optimization of  $\pi_z$  and  $q_\lambda(z|s)$  and discovery of meaningful behaviors difficult.

Furthermore, spectral normalization can improve the performance of variational empowerment algorithms in more challenging control tasks, as will be discussed in Section 6.3. These results confirm that *high expressivity does not necessarily mean better performance*, and therefore that inductive biases or *proper regularizations on the posterior*  $q_\lambda(z|s)$  are important for the performance of the algorithm and the quality of the goal representation learned.

2D Pointmass	SN?	$\mathcal{F}$	LGR( $z$ )
$ G  = 10$	- ✓	-0.38 <b>-0.14</b>	0.94 <b>0.96</b>
$ G  = 20$	- ✓	-0.86 <b>-0.24</b>	0.91 <b>0.96</b>
$ G  = 50$	- ✓	<b>-1.02</b> -1.15	<b>0.83</b> 0.78

Table 2: DIAYN v.s. DIAYN + Spectral Normalization (SN) on 2D PointMass (See Figure 3). Please refer to Section 6.3 for details of the metrics and discussions.

## 6 Variational Empowerment as Goal-Conditioned RL

We have seen that goal-conditioned RL can be viewed as a special case of the unified optimization objective in Eq. (2). A natural question that follows is whether training methods that work well for GCRL can also be used within the more general VGCRRL framework. We answer in the affirmative, deriving an extension to a goal-relabeling technique (HER) and an evaluation metric for variational empowerment algorithms. We present some techniques that are helpful for learning VGCRRL.

### 6.1 P-HER: Posterior Hindsight Experience Replay

As discussed in Section 3.2, Hindsight Experience Replay (HER) (Andrychowicz et al., 2017) substantially improves off-policy learning of goal-conditioned policies and value functions. Inspired by the mathematical connection to GCRL, we can derive an equivalent of HER or goal relabeling in the context of VGCRRL. Specifically, we relabel  $z$  for fitting the policy  $\pi_\theta(a|s, z)$  with the relabeled goal  $\mathbb{S}(s_{0:T})$  derived from the final state  $s_T$ :

$$\mathbb{S}(s_{0:T}) \sim q_\lambda(z | s_T). \quad (7)$$

We call this relabeling technique *Posterior HER (P-HER)*, as it can be seen as an application of posterior sampling (Hausman et al., 2018; Rakelly et al., 2019) to HER. Since  $q_\lambda$  is non-stationary, we always use the *up-to-date* estimate of  $q_\lambda(z|s_T)$  with the latest parameter  $\lambda$  of posterior  $q_\lambda$ . This is the biggest difference to the vanilla HER where the goal mapping function is fixed. We note that other state sampling distributions of HER, e.g., uniform over  $s_{0:T}$  or  $s_{t>k}$  can also be trivially supported. For fitting the discriminator  $q_\lambda(z|s)$ , we do not relabel the goal. Similarly to HER in GCRL, this technique can be viewed as a curriculum for enabling parameterized policy optimization to focus on high-reward regions. This technique can *accelerate optimization steps* for  $\pi_\theta$  (Equation 4) in the same spirit of HER when latent goal reaching is non-trivial, especially in high-dimensional state spaces and with immaturely-shaped reward functions.

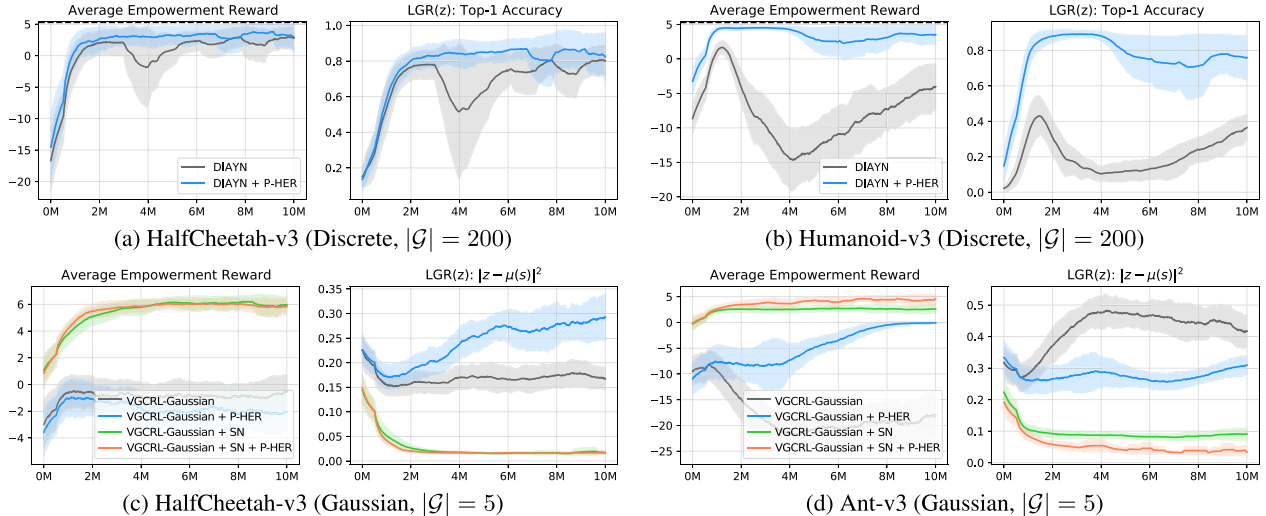


Figure 4: Learning curves (selected) of variational empowerment methods with *discrete* ((a)-(b), see Table 3) and *continuous* ((c)-(d), see Table 4) goal spaces. We can see that Posterior HER (P-HER) makes optimization of the MI objective faster and more stable, especially when the goal space is large and/or goal-conditioned policy is difficult to learn. The use of spectral normalization (SN, Section 5.3) provides performance gains for VGCRl with Gaussian discriminators. The higher  $LGR(z)$  is, the better in discrete cases (accuracy; (a)-(b)); the lower  $LGR(z)$  is, the better in continuous cases (distance; (c)-(d)). Learning curves are averaged over 3 random seeds. More plots can be found in Appendix C.

## 6.2 Latent Goal Reaching: A Metric for MI-Based Empowerment Algorithms

One limitation in mutual information-based RL (Eysenbach et al., 2019; Sharma et al., 2020b) is lack of objective evaluation metrics. Previous works directly evaluate qualitative results of learned behaviors, but quantitative metrics are limited to discriminator rewards (i.e.,  $\mathbb{E}[\log q(z|s) - p(z)]$ ) which can saturate and do not always correspond to the quality of learned behaviors, or downstream task evaluations such as exploration bonus and pre-trained primitives for hierarchical RL (Florensa et al., 2017; Eysenbach et al., 2019; Sharma et al., 2020b; Hansen et al., 2020).

By contrast, goal-conditioned RL has a clear metric: define a sample of goals and evaluate average goal-reaching performance. Since we demonstrated how MI-based RL is closely related to GCRL, we propose Latent Goal Reaching (LGR) as a new metric for evaluating MI-based algorithms such as DIAYN or DADS. The procedure is described in Algorithm 1. *It allows us to evaluate MI-based RL as just another goal-reaching problem:* we measure how accurately the goal-conditioned policy can reach the given goal states of interest. This metric measures the quality of both the discriminator (i.e., goal representation) and the goal-conditioned policy. We note the range of this metric is not dependent on the choice of distribution family of  $q(\cdot)$ , allowing comparison across different types of  $q(z|s)$ .

The value of  $LGR(s)$  metric depends on the choice of target states. In locomotion control tasks (Brockman et al.,

---

### Algorithm 1 Latent Goal Reaching Metric: $LGR(s)$

---

**Input:** Target states  $s^{1:N}$ , trained  $\pi_\theta(a|s, z)$ ,  $q_\lambda(z|s)$

**Output:** Average distance  $\bar{d}$  to goal states

**for**  $i \leftarrow 1$  **to**  $N$  **do**

Embed target state into a goal:  $z^i \leftarrow \mathbb{E}[q_\lambda(\cdot|s^i)]$   
 Run  $\pi(\cdot|s^i, z^i)$  for  $T$  time steps, observe final state  $s_T^i$   
 Compute  $d(s^i, s_T^i)$  (e.g., squared distance)

**end**

Report the average over  $N$  episodes:  $\bar{d} = \sum_i d(s^i, s_T^i)/N$

---

2016), we would often want to discover and learn behaviors where the robot can walk or move (Eysenbach et al., 2019; Sharma et al., 2020b). To evaluate this, one can generate diverse target states of moving in different directions and at different velocities from expert policies. In such cases we can compute the distance  $d(s^i, s_T^i)$  between target and achieved states with respect to velocity dimensions (e.g., velocity in  $x$  and  $y$  axis). We call this variant of the LGR metric as  $LGR_v(s)$ . Details of target state generation used in our experiments are given in Appendix D.

We also consider the  $LGR(z)$  metric which measures the goal reaching performance in the latent goal space. For discrete  $z$ , this simply can be the top-1 accuracy of the discriminator  $q(z|s)$  with respect to the true goal  $z \sim p(z)$ . For continuous  $z$ ,  $LGR(z)$  is the squared distance  $\|z - \operatorname{argmax} q(z|s)\|^2$  between the goal and the mode of the discriminator  $\operatorname{argmax} q(z|s)$ . We note that this metric agrees the state distance metric used in Laplacian embedding (Wu et al., 2018), i.e.,  $\|\phi(s) - \phi(h^{-1}(z))\|$ , with a

$ \mathcal{G} $	Method	HalfCheetah			Ant			Humanoid		
		$\mathcal{F}$	$\text{LGR}_v(s)$	$\text{LGR}(z)$	$\mathcal{F}$	$\text{LGR}_v(s)$	$\text{LGR}(z)$	$\mathcal{F}$	$\text{LGR}_v(s)$	$\text{LGR}(z)$
10	DIAYN	<b>1.608</b>	<b>0.800</b>	<b>0.963</b>	-0.835	0.529	0.806	1.261	0.523	0.922
	DIAYN + P-HER	1.372	1.424	0.934	<b>-0.049</b>	<b>0.486</b>	<b>0.889</b>	<b>1.856</b>	<b>0.312</b>	<b>0.953</b>
20	DIAYN	1.732	<b>1.125</b>	<b>0.920</b>	-1.308	0.610	0.763	0.713	0.315	0.768
	DIAYN + P-HER	<b>1.852</b>	1.214	0.891	<b>-1.288</b>	<b>0.515</b>	<b>0.823</b>	<b>2.251</b>	<b>0.183</b>	<b>0.922</b>
50	DIAYN	1.475	<b>0.673</b>	0.827	-3.812	<b>0.402</b>	0.523	-1.158	<b>0.268</b>	0.549
	DIAYN + P-HER	<b>1.699</b>	0.704	<b>0.834</b>	<b>-2.171</b>	0.637	<b>0.750</b>	<b>2.848</b>	0.545	<b>0.891</b>
200	DIAYN	2.854	<b>0.698</b>	0.801	-8.450	<b>0.396</b>	0.113	-4.396	<b>0.208</b>	0.337
	DIAYN + P-HER	<b>3.357</b>	0.814	<b>0.844</b>	<b>-7.047</b>	0.555	<b>0.263</b>	<b>3.448</b>	0.866	<b>0.766</b>
1000	DIAYN	-8.286	1.156	0.176	-16.795	0.434	0.005	-8.914	<b>0.325</b>	0.101
	DIAYN + P-HER	<b>-4.424</b>	<b>0.510</b>	<b>0.361</b>	<b>-10.941</b>	<b>0.322</b>	<b>0.028</b>	<b>3.395</b>	1.569	<b>0.762</b>

Table 3: Evaluation of Latent Goal-Reaching Metric on MuJoCo control suites, after a total of 10M environment steps of training.  $\mathcal{F}$  is the (average) empowerment reward,  $\mathcal{F} = \mathbb{E}_z[\log q_\lambda(z|s) - p(z)]$ .  $\text{LGR}_v(s)$  is the squared error in observation space (the lower, the better) with respect to velocity dimensions between the marginal state and the target state, and  $\text{LGR}(z)$  denotes the accuracy of top-1 classification of the discriminator  $q_\lambda(z|s)$  (the higher, the better, max 1.0).

$ \mathcal{G}  = 5$	$q_\lambda(z s)$	P-HER?	SN?	HalfCheetah			Ant			Humanoid		
				$\mathcal{F}$	$\text{LGR}_v(s)$	$\text{LGR}(z)$	$\mathcal{F}$	$\text{LGR}_v(s)$	$\text{LGR}(z)$	$\mathcal{F}$	$\text{LGR}_v(s)$	$\text{LGR}(z)$
	$\mathcal{N}(\mu(s), \text{fixed}^2)$	-	-	0.932	1.005	0.159	-0.590	1.005	0.382	0.239	1.461	0.202
		✓	-	-0.142	1.273	0.360	0.140	2.449	0.300	0.020	1.452	0.244
	$\mathcal{N}(\mu(s), \Sigma(s)^2)$	-	-	-0.731	1.251	0.172	-18.490	<b>0.306</b>	0.427	-3.597	0.538	0.147
		✓	-	-2.161	1.132	0.289	-0.108	2.423	0.303	1.207	0.206	0.074
		-	✓	<b>5.856</b>	<b>0.604</b>	0.019	2.548	0.925	0.091	4.509	0.460	0.040
		✓	✓	5.803	1.352	<b>0.017</b>	<b>4.349</b>	0.463	<b>0.039</b>	<b>5.203</b>	<b>0.203</b>	<b>0.026</b>
	GMM ( $K = 8$ )	-	-	-2.646	0.766	0.325	-16.196	0.367	0.486	-3.576	0.231	0.198
		✓	-	-3.091	1.065	0.404	-2.874	2.794	0.325	3.526	0.581	0.043

Table 4: Comparison of continuous variants of VGCRl, where the dimension of the goal space is  $|\mathcal{G}| = 5$ . SN denotes Spectral Normalization (Section 5.3).  $\text{LGR}(z)$  is the goal reaching performance in the latent space (the lower, the better). A full table containing more comprehensive comparison and corresponding plots can be found in Appendix C (Table 5).

state embedding  $\phi(s) := \operatorname{argmax}_z q(z|s)$ , where  $h^{-1}(z)$  is defined to be an arbitrary state  $s$  that is associated with latent  $z$ , which is in our case the marginal state from the latent-conditioned policy  $\pi_\theta$ . A difference is that (Wu et al., 2018) use contrastive learning whereas VGCRl maximizes likelihood to learn the representation  $q$ .

### 6.3 Experiments: Posterior HER and Latent Goal Reaching

In this section, we evaluate the performance of several variants of VGCRl on standard locomotion tasks (Brockman et al., 2016). We consider both a discrete latent space, analogous to that used by DIAYN (Eysenbach et al., 2019), and a continuous latent space, where the variational posterior  $q_\lambda(z|s)$  chosen to be a Gaussian posterior (VGCRl-Gaussian), with either learnable or fixed variances, or Gaussian Mixtures (VGCRl-GMM). To evaluate the performance, we report the following metrics: (1) the empowerment objective  $\mathcal{F} = \mathbb{E}_z[\log q(z|s) - p(z)]$ , (2)  $\text{LGR}(z)$ : the latent goal reaching metric, and (3)  $\text{LGR}_v(s)$ : the latent goal reaching metric with respect to velocity dimensions (Section 6.2).

We first observe that P-HER can accelerate and improve learning (Figure 4, Tables 3 and 4). Such improvements are significant in high-dimensional goal spaces (e.g.,  $|\mathcal{G}| = 200$ ) and more difficult control tasks such as Ant or Humanoid with high dimensionalities, in which both the discriminator and the latent-conditioned policy are difficult to learn. Because an optimization of goal-conditioned policy (Eq. 4) is more difficult than discriminators (Eq. 3), relabeling of goal can greatly accelerate RL, which also results in better discriminability. As shown in Tables 4 and 5, P-HER can improve not only the optimization objective but also other metrics such as discriminator’s accuracy and goal reaching performance across different design choices and environments.

Moreover, when Spectral Normalization (Section 5.3) is applied, we observe a significant improvement in terms of the learning progress and evaluation metrics (Figure 4, Table 4). As shown in Figure 4, while there are some progress with vanilla VGCRl-Gaussian (or VGCRl-GMM), the optimization objective  $\mathcal{F}$  as well as evaluation metrics do not improve as much as the one with SN, despite the bigger



expressivity due to no constraint on  $q_\lambda(z|s)$ . We can also see that with SN (and P-HER as well), the distance between achieved and desired goal is much lower. More experimental results and discussions can be found in Appendix C.

## 7 Conclusion

Our variational GCRL (VGCR) framework unifies unsupervised skill learning methods based on variational empowerment (Eysenbach et al., 2019; Gregor et al., 2017; Sharma et al., 2020b) with goal-conditioned RL (GCRL) methods, allowing us to transfer techniques and insights across both types of approaches. Viewing GCRL as variational empowerment, we derive simple extensions of goal-based methods that exhibit some representation learning capability of variational methods, e.g., disentangle underlying factors of variations and automatically determine controllable dimensions, while keeping the learning stability of GCRL. Viewing variational empowerment as GCRL, we can transfer popular optimization techniques such as relabeling from GCRL to variational empowerment algorithms, and propose latent goal-reaching (LGR) as a more objective, performance-based metric for evaluating the quality of skill (latent goal) discovery. We hope that these insights can lay the ground for further developments of more capable and performant algorithms for unsupervised reinforcement learning in future work.

## Acknowledgements

The authors are grateful to Igor Mordatch, Lisa Lee, Aleksandra Faust, and Nicolas Heess for helpful discussions and comments. JC was partly supported by Korea Foundation for Advanced Studies.

## References

- Abbeel, P. and Ng, A. Y. Apprenticeship learning via inverse reinforcement learning. In *ICML*, 2004.
- Achiam, J., Edwards, H., Amodei, D., and Abbeel, P. Variational Option Discovery Algorithms. *arXiv preprint arXiv:1807.10299*, 2018.
- Andrychowicz, M., Wolski, F., Ray, A., Schneider, J., Fong, R., Welinder, P., McGrew, B., Tobin, J., Abbeel, P., and Zaremba, W. Hindsight Experience Replay. In *NIPS*, 2017.
- Barber, D. and Agakov, F. V. The IM algorithm: a variational approach to information maximization. In *NIPS*, 2003.
- Barreto, A., Dabney, W., Munos, R., Hunt, J. J., Schaul, T., Silver, D., and van Hasselt, H. Successor Features for Transfer in Reinforcement Learning. In *NIPS*, pp. 1–24, November 2017.
- Bellemare, M., Srinivasan, S., Ostrovski, G., Schaul, T., Saxton, D., and Munos, R. Unifying count-based exploration and intrinsic motivation. In *Advances in neural information processing systems*, pp. 1471–1479, 2016.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- Campos, V., Trott, A., Xiong, C., Socher, R., i Nieto, X. G., and Torres, J. Explore, discover and learn: Unsupervised discovery of state-covering skills, 2020.
- Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., and Amodei, D. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems*, pp. 4299–4307, 2017.
- Ciosek, K. and Whiteson, S. Expected policy gradients. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Colas, C., Fournier, P., Sigaud, O., Chetouani, M., and Oudeyer, P.-Y. Curious: intrinsically motivated modular multi-goal reinforcement learning. *arXiv preprint arXiv:1810.06284*, 2018.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Eysenbach, B., Gupta, A., Ibarz, J., and Levine, S. Diversity is All You Need: Learning Skills without a Reward Function. In *ICLR*, 2019.
- Florensa, C., Duan, Y., and Abbeel, P. Stochastic neural networks for hierarchical reinforcement learning. *arXiv preprint arXiv:1704.03012*, 2017.
- Fox, R., Pakman, A., and Tishby, N. Taming the noise in reinforcement learning via soft updates. *arXiv preprint arXiv:1512.08562*, 2015.
- Fu, J., Luo, K., and Levine, S. Learning robust rewards with adversarial inverse reinforcement learning. *arXiv preprint arXiv:1710.11248*, 2017.
- Ghasemipour, S. K. S., Zemel, R., and Gu, S. A divergence minimization perspective on imitation learning methods. *Conference on Robot Learning (CoRL)*, 2019.
- Gregor, K., Rezende, D. J., and Wierstra, D. Variational Intrinsic Control. In *ICLR Workshop*, 2017.
- Gu, S., Holly, E., Lillicrap, T., and Levine, S. Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In *2017 IEEE international conference on robotics and automation (ICRA)*, pp. 3389–3396. IEEE, 2017a.

- Gu, S. S., Lillicrap, T., Turner, R. E., Ghahramani, Z., Schölkopf, B., and Levine, S. Interpolated policy gradient: Merging on-policy and off-policy gradient estimation for deep reinforcement learning. In *Advances in neural information processing systems*, pp. 3846–3855, 2017b.
- Guadarrama, S., Korattikara, A., Ramirez, O., Castro, P., Holly, E., Fishman, S., Wang, K., Gonina, E., Wu, N., Kokiopoulou, E., Sbaiz, L., Smith, J., Bartók, G., Berent, J., Harris, C., Vanhoucke, V., and Brevdo, E. TF-Agents: A library for reinforcement learning in tensorflow. 2018. URL <https://github.com/tensorflow/agents>.
- Gupta, A., Eysenbach, B., Finn, C., and Levine, S. Unsupervised meta-learning for reinforcement learning. *arXiv preprint arXiv:1806.04640*, 2018.
- 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. *arXiv preprint arXiv:1812.05905*, 2018.
- Hadfield-Menell, D., Milli, S., Abbeel, P., Russell, S. J., and Dragan, A. Inverse reward design. In *Advances in neural information processing systems*, pp. 6765–6774, 2017.
- Hansen, S., Dabney, W., Barreto, A., Van de Wiele, T., Warde-Farley, D., and Mnih, V. Fast Task Inference with Variational Intrinsic Successor Features. In *ICLR*, 2020.
- Hausman, K., Springenberg, J. T., Wang, Z., Heess, N., and Riedmiller, M. Learning an embedding space for transferable robot skills. 2018.
- Hazan, E., Kakade, S. M., Singh, K., and Van Soest, A. Provably efficient maximum entropy exploration. *arXiv preprint arXiv:1812.02690*, 2018.
- Ho, J. and Ermon, S. Generative adversarial imitation learning. In *Advances in neural information processing systems*, pp. 4565–4573, 2016.
- Jung, T., Polani, D., and Stone, P. Empowerment for continuous agent—environment systems. *Adaptive Behavior*, 19(1):16–39, 2011.
- Kaelbling, L. P. Learning to achieve goals. In *IJCAI*, pp. 1094–1099. Citeseer, 1993.
- Kakade, S. M. A natural policy gradient. In *Advances in neural information processing systems*, pp. 1531–1538, 2002.
- Kalashnikov, D., Irpan, A., Pastor, P., Ibarz, J., Herzog, A., Jang, E., Quillen, D., Holly, E., Kalakrishnan, M., Vanhoucke, V., et al. Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation. *arXiv preprint arXiv:1806.10293*, 2018.
- Karl, M., Soelch, M., Becker-Ehmck, P., Benbouzid, D., van der Smagt, P., and Bayer, J. Unsupervised real-time control through variational empowerment, 2017.
- Klyubin, A. S., Polani, D., and Nehaniv, C. L. All else being equal be empowered. In *European Conference on Artificial Life*, pp. 744–753. Springer, 2005.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- Lee, L., Eysenbach, B., Parisotto, E., Xing, E., Levine, S., and Salakhutdinov, R. Efficient exploration via state marginal matching. *arXiv preprint arXiv:1906.05274*, 2019.
- Li, Y., Song, J., and Ermon, S. Infogail: Interpretable imitation learning from visual demonstrations. In *Advances in neural information processing systems*, 2017.
- Lynch, C., Khansari, M., Xiao, T., Kumar, V., Tompson, J., Levine, S., and Sermanet, P. Learning latent plans from play. *arXiv preprint arXiv:1903.01973*, 2019.
- Maddison, C. J., Lawson, D., Tucker, G., Heess, N., Doucet, A., Mnih, A., and Teh, Y. W. Particle value functions. *arXiv preprint arXiv:1703.05820*, 2017.
- Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Mohamed, S. and Rezende, D. J. Variational information maximisation for intrinsically motivated reinforcement learning. In *Advances in neural information processing systems*, pp. 2125–2133, 2015.
- Nachum, O., Gu, S., Lee, H., and Levine, S. Near-optimal representation learning for hierarchical reinforcement learning. *arXiv preprint arXiv:1810.01257*, 2018.
- Nair, A. V., Pong, V., Dalal, M., Bahl, S., Lin, S., and Levine, S. Visual reinforcement learning with imagined goals. In *Advances in Neural Information Processing Systems*, pp. 9191–9200, 2018.
- Ng, A. Y., Russell, S. J., et al. Algorithms for inverse reinforcement learning. In *ICML*, 2000.
- Oudeyer, P.-Y. and Kaplan, F. What is intrinsic motivation? a typology of computational approaches. *Frontiers in neurobotics*, 1:6, 2009.

- Pathak, D., Agrawal, P., Efros, A. A., and Darrell, T. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 16–17, 2017.
- Plappert, M., Andrychowicz, M., Ray, A., McGrew, B., Baker, B., Powell, G., Schneider, J., Tobin, J., Chociej, M., Welinder, P., et al. Multi-goal reinforcement learning: Challenging robotics environments and request for research. *arXiv preprint arXiv:1802.09464*, 2018.
- Pong, V., Gu, S., Dalal, M., and Levine, S. Temporal difference models: Model-free deep rl for model-based control. *International Conference on Learning Representations (ICLR)*, 2018.
- Pong, V. H., Dalal, M., Lin, S., Nair, A., Bahl, S., and Levine, S. Skew-fit: State-covering self-supervised reinforcement learning. *arXiv preprint arXiv:1903.03698*, 2019.
- Rakelly, K., Zhou, A., Quillen, D., Finn, C., and Levine, S. Efficient off-policy meta-reinforcement learning via probabilistic context variables. *arXiv preprint arXiv:1903.08254*, 2019.
- Rudner, T. G. J., Pong, V. H., McAllister, R., Gal, Y., and Levine, S. Outcome-driven reinforcement learning via variational inference. 2021.
- Salge, C., Glackin, C., and Polani, D. Empowerment – An Introduction. In *Guided Self-Organization: Inception*. 2014.
- Schaul, T., Horgan, D., Gregor, K., and Silver, D. Universal Value Function Approximators. In *ICML*, 2015.
- Schmidhuber, J. Curious model-building control systems. In *Proc. international joint conference on neural networks*, pp. 1458–1463, 1991.
- Schmidhuber, J. Formal theory of creativity, fun, and intrinsic motivation (1990–2010). *IEEE Transactions on Autonomous Mental Development*, 2(3):230–247, 2010.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. Trust region policy optimization. In *International conference on machine learning*, pp. 1889–1897, 2015.
- Sharma, A., Ahn, M., Levine, S., Kumar, V., Hausman, K., and Gu, S. Emergent real-world robotic skills via unsupervised off-policy reinforcement learning. In *Robotics: Science and Systems (RSS)*, 2020a.
- Sharma, A., Gu, S., Levine, S., Kumar, V., and Hausman, K. Dynamics-aware unsupervised discovery of skills. In *ICLR*, 2020b.
- Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D., and Riedmiller, M. Deterministic policy gradient algorithms. 2014.
- Sutton, R. S., McAllester, D. A., Singh, S. P., and Mansour, Y. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pp. 1057–1063, 2000.
- Tassa, Y., Doron, Y., Muldal, A., Erez, T., Li, Y., Casas, D. d. L., Budden, D., Abdolmaleki, A., Merel, J., Lefrancq, A., et al. Deepmind control suite. *arXiv preprint arXiv:1801.00690*, 2018.
- Todorov, E., Erez, T., and Tassa, Y. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033. IEEE, 2012.
- Warde-Farley, D., Van de Wiele, T., Kulkarni, T., Ionescu, C., Hansen, S., and Mnih, V. Unsupervised Control Through Non-Parametric Discriminative Rewards. In *ICLR*, 2019.
- Watkins, C. J. and Dayan, P. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.
- Wipf, D. P. and Nagarajan, S. S. A new view of automatic relevance determination. In *Advances in neural information processing systems*, pp. 1625–1632, 2008.
- Wu, Y., Tucker, G., and Nachum, O. The laplacian in rl: Learning representations with efficient approximations. *arXiv preprint arXiv:1810.04586*, 2018.
- Ziebart, B. D., Maas, A. L., Bagnell, J. A., and Dey, A. K. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, pp. 1433–1438. Chicago, IL, USA, 2008.