

Off-Policy Deep Reinforcement Learning with Analogous Disentangled Exploration

Anji Liu

University of California, Los Angeles
anjiliu219@gmail.com

Yitao Liang

University of California, Los Angeles
yliang@cs.ucla.edu

Guy Van den Broeck

University of California, Los Angeles
guyvdb@cs.ucla.edu

ABSTRACT

Off-policy reinforcement learning (RL) is concerned with learning a rewarding policy by executing another policy that gathers samples of experience. While the former policy (i.e. target policy) is rewarding but in-expressive (in most cases, deterministic), doing well in the latter task, in contrast, requires an expressive policy (i.e. behavior policy) that offers guided and effective exploration. Contrary to most methods that make a trade-off between optimality and expressiveness, disentangled frameworks explicitly decouple the two objectives, which each is dealt with by a distinct separate policy. Although being able to freely design and optimize the two policies with respect to their own objectives, naively disentangling them can lead to inefficient learning or stability issues. To mitigate this problem, our proposed method *Analogous Disentangled Actor-Critic* (ADAC) designs analogous pairs of actors and critics. Specifically, ADAC leverages a key property about Stein variational gradient descent (SVGD) to constraint the expressive energy-based behavior policy with respect to the target one for effective exploration. Additionally, an analogous critic pair is introduced to incorporate intrinsic rewards in a principled manner, with theoretical guarantees on the overall learning stability and effectiveness. We empirically evaluate environment-reward-only ADAC on 14 continuous-control tasks and report the state-of-the-art on 10 of them. We further demonstrate ADAC, when paired with intrinsic rewards, outperform alternatives in exploration-challenging tasks.

KEYWORDS

Reinforcement Learning; Deep Reinforcement Learning; Exploration

ACM Reference Format:

Anji Liu, Yitao Liang, and Guy Van den Broeck. 2020. Off-Policy Deep Reinforcement Learning with Analogous Disentangled Exploration. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*, Auckland, New Zealand, May 9–13, 2020, IFAAMAS, 9 pages.

1 INTRODUCTION

Reinforcement learning (RL) studies the control problem where an agent tries to navigate through an unknown environment [35]. The agent attempts to maximize its cumulative rewards through an iterative trial-and-error learning process [1]. Recently, we have seen many successes of applying RL to challenging simulation [20, 25] and real-world [19, 34, 39] problems. Inherently, RL consists of two distinct but closely related objectives: learn the best possible

policy from the gathered samples (i.e. *exploitation*) and collect new samples effectively (i.e. *exploration*). While the *exploitation* step shares certain similarities with tasks such as supervised learning, *exploration* is unique, essential, and is often viewed as the backbone of many successful RL algorithms [15, 24].

In order to explore novel states that are potentially rewarding, it is crucial to incorporate randomness when interacting with the environment. Thanks to its simplicity, injecting noise into the action [11, 21] or parameter space [9, 29] is widely used to implicitly construct behavior policies from target policies. In most prior work, the injected noise has a mean of zero, such that the updates to the target policy have no bias [12, 13]. The stability of noise-based exploration, which is obtained from its non-biased nature, makes it a safe exploration strategy. However, noise-based approaches are generally less effective since they are neither aware of potentially rewarding actions nor guided by the exploration-oriented targets.

To tackle the above problem, two orthogonal lines of approaches have been proposed. One of them considers extracting more information from the current knowledge (i.e. gathered samples). For example, energy-based RL algorithms learn to capture potentially rewarding actions through its energy objective [15, 35]. A second line of work considers leveraging external guidance to aid exploration. In a nutshell, they formulate some intuitive tendencies in exploration as an additional reward function called intrinsic reward [2, 17]. Guided by these auxiliary tasks, RL algorithms tend to act curiously, substantially improving exploration of the state space.

Despite their promising exploration efficiency, both lines of work fail to fully *exploit* the collected samples and turn them into the highest performing policy, as their learned policy often executes sub-optimal actions. To avoid this undesirable *exploration-exploitation* trade-off, several attempts have been made to separately design two policies (i.e. disentangle them), of which one aims to gather the most informative examples (and hence is commonly referred as the *behavior policy*) while the other attempts to best utilize the current knowledge from the gathered samples (and hence is usually referred as the *target policy*) [4, 6]. To help fulfill their respective goals, disentangled objective functions and learning paradigms are further designed and separately applied to the two policies.

However, naively disentangling the behavior from the target policy would render their update process unstable. For example, when disentangled naively, the two policies tend to differ substantially due to their contrasting objectives, which is known to potentially result in catastrophic learning failure [27]. To mitigate this problem, we propose *Analogous Disentangled Actor-Critic* (ADAC), where being *analogous* is reflected by the constraints imposed on the disentangled actor-critic [23] pairs. ADAC consists of two main algorithmic contributions. First, *policy co-training* guides the behavior policy's update by the target policy, making the gathered

samples more helpful for the target policy’s learning process while keeping the expressiveness of the behavior policy for extensive exploration (Section 4.2). Second, *critic bounding* allows an additional explorative critic to be trained with the aid of intrinsic rewards (Section 4.3). Under certain constraints from the target policy, the resultant critic maintains the curiosity incentivized by intrinsic rewards while guarantees training stability of the target policy.

Besides Section 4’s elaboration of our method, the rest of the paper is organized as follows. Section 2 reviews and summarizes the related work. Key background concepts and notations are introduced in Section 3. Experiment details of ADAC are explained in Section 5. Finally, conclusions are presented in Section 6.¹

2 RELATED WORK

Learning to be aware of potentially rewarding actions is a promising strategy to conduct exploration, as it automatically prunes less rewarding actions and concentrates exploration efforts on those with high potential. To capture these actions, expressive learning models/objectives are widely used. Most noticeable recent work on this direction, such as Soft Actor-Critic [15], EntRL [31], and Soft Q Learning [14], learns an expressive energy-based target policy according to the maximum entropy RL objective [43]. However, the expressiveness of their policies in turn becomes a burden for their optimality, and in practice, trade-offs such as temperature controlling [16] and reward scaling [14] have to be made for better overall performance. As we shall show later, ADAC makes use of a similar but extended energy-based target, and alleviates the compromise on optimality using the *analogous disentangled* framework.

Ad-hoc exploration-oriented learning targets that are designed to better explore the state space are also promising. Some recent research efforts on this line include count-based exploration [2, 42] and intrinsic motivation [10, 17, 18] approaches. The outcome of these methods is usually an auxiliary reward termed the *intrinsic reward*, which is extremely useful when the environment-defined reward is sparsely available. However, as we shall illustrate in Section 5.3, intrinsic reward potentially biases the task-defined learning objective, leading to catastrophic failure in some tasks. Again, with the disentangled nature of ADAC, we give a principled solution to solve this problem with theoretical guarantees (Section 4.3).

Explicitly disentangling exploration from exploitation has been used to solve a common problem in the above approaches, which is sacrificing the target policy’s optimality for better exploration. By separately designing exploration and exploitation components, both objectives can be better pursued simultaneously. Specifically, GEP-PG [6] uses a Goal Exploration Process (GEP) [8] to generate samples and feed them to the replay buffer of DDPG [21] or its variants. Multiple losses for exploration (MULEX) [4] proposes to use a series of intrinsic rewards to optimize different policies in parallel, which in turn generates abundant samples to train the target policy. Despite having intriguing conceptual ideas, they overlook the training stability issue caused by the mismatch in the distribution of collected samples (using the behavior policy) and the distribution induced by the target policy, which is formalized as *extrapolation*

error in [12]. ADAC aims to mitigate the training stability issue caused by the extrapolation error while maintaining effective exploration exploitation trade-off promised by expressive behavior policies (Section 4.2) as well as intrinsic rewards (Section 4.3) using its *analogous disentangled* actor-critic pairs.

3 PRELIMINARIES

In this section, we introduce the RL setting we address in this paper, and some background concepts that we utilize to build our method.

3.1 RL with Continuous Control

In a standard *reinforcement learning* (RL) setup, an agent interacts with an unknown environment at discrete time steps and aims to maximize the *reward* signal [35]. The environment is often formalized as a *Markov Decision Process* (MDP), which can be succinctly defined as a 5-tuple $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{P}, \gamma \rangle$. At time step t , the agent in state $s_t \in \mathcal{S}$ takes action $a_t \in \mathcal{A}$ according to *policy* π , a conditional distribution of a given s , leading to the next state s_{t+1} according to the transition probability $\mathcal{P}(s_{t+1} | s_t, a_t)$. Meanwhile, the agent observes reward $r_t \sim \mathcal{R}(s_t, a_t)$ emitted from the environment.²

The agent strives to learn the *optimal policy* that maximizes the expected return $J(\pi) = \mathbb{E}_{s_0 \sim \rho_0, a_t \sim \pi, s_{t+1} \sim \mathcal{P}, r_t \sim \mathcal{R}} [\sum_{t=0}^{\infty} \gamma^t r_t]$, where ρ_0 is the initial state distribution and $\gamma \in [0, 1)$ is the discount factor balancing the priority of short and long-term rewards. For continuous control, the policy π (also known as the actor in the actor-critic framework) parameterized by θ can be updated by taking the gradient $\nabla_{\theta} J(\pi)$. According to the deterministic policy gradient theorem [33], $\nabla_{\theta} J(\pi) = \mathbb{E}_{(s,a) \sim \rho_{\pi}} [\nabla_a Q_{\mathcal{R}}^{\pi}(s, a) \nabla_{\theta} \pi(s)]$, where ρ_{π} denotes the state-action marginals of the trajectory distribution induced by π , and $Q_{\mathcal{R}}^{\pi}$ denotes the state-action value function (also known as the critic in the actor-critic framework), which represents the expected return under the reward function specified by \mathcal{R} when performing action a at state s and following policy π afterwards. Intuitively, it measures how preferable executing action a is at state s with respect to the policy π and reward function \mathcal{R} . Following [3], we additionally introduce the Bellman operator, which is commonly used to update the Q -function. The Bellman operator $\mathcal{T}_{\mathcal{R}}^{\pi}$ uses \mathcal{R} and π to update an arbitrary value function Q , which is not necessarily defined with respect to the same π or \mathcal{R} . For example, the outcome of $\mathcal{T}_{\mathcal{R}_1}^{\pi_1} Q_{\mathcal{R}_2}^{\pi_2}(s_t, a_t)$ is defined as $\mathcal{R}_1(s_t, a_t) + \gamma \mathbb{E}_{s_{t+1} \sim \mathcal{P}, a_{t+1} \sim \pi_1} [Q_{\mathcal{R}_2}^{\pi_2}(s_{t+1}, a_{t+1})]$. By slightly abusing notations, we further define the outcome of $\mathcal{T}_{\mathcal{R}_1}^{\max} Q_{\mathcal{R}_2}^{\pi_2}(s_t, a_t)$ as $\mathcal{R}_1(s_t, a_t) + \gamma \max_{a_{t+1}} \mathbb{E}_{s_{t+1} \sim \mathcal{P}} [Q_{\mathcal{R}_2}^{\pi_2}(s_{t+1}, a_{t+1})]$. Some also call $\mathcal{T}_{\mathcal{R}}^{\max}$ the Bellman optimality operator.

3.2 Off-policy Learning and Behavior Policy

To aid exploration, it is a common practice to construct/store more than one policy for the agent (either implicitly or explicitly). Off-policy actor-critic methods [40] allow us to make a clear separation between the *target policy*, which refers to the best policy currently learned by the agent, and the *behavior policy*, which the agent follows to interact with the environment. Note that the discussion in Section 3.1 is largely around the target policy. Thus, starting from this point, to avoid confusion, π is reserved to only denote the target

¹A longer version of this paper with supplementary material is available at <https://arxiv.org/abs/2002.10738>; code to reproduce our experiments at <https://github.com/UCLA-StarAI/Analogous-Disentangled-Actor-Critic>.

²In all the environments considered in this paper, actions are assumed to be continuous.




Figure 1: Evaluation of the amortized SVGD learning algorithm [7] (Eq (3)) with different β under two target distributions.

policy and notation μ is introduced to denote the behavior policy. Due to the policy separation, the target policy π is instead resorting to the estimates calculated with regards to samples collected by the behavior policy μ , that is, the deterministic policy gradient mentioned above is approximated as

$$\nabla_{\theta} J_{\pi}(\theta) \approx \mathbb{E}_{(s,a) \sim \rho_{\mu}} [\nabla_a Q_{\mathcal{R}}^{\pi}(s,a) \nabla_{\theta} \pi(s)], \quad (1)$$

where \mathcal{R} is the environment-defined reward. One of the most notable off-policy learning algorithms that capitalize on this idea is deep deterministic policy gradient (DDPG) [21]. To mitigate function approximation errors in DDPG, Fujimoto et al. proposes TD3 [11]. Given that DDPG and TD3 have demonstrated themselves to be competitive in many continuous control benchmarks, we choose to implement our *Analogous Disentangled Actor Critic* (ADAC) on top of their target policies. Yet, it is worth reiterating that ADAC is compatible with any existing off-policy learning algorithms. We defer a more detailed discussion of ADAC’s compatibility until we start formally introducing our method in Section 4.1.

3.3 Expressive Behavior Policies through Energy-Based Representation

One promising way to design an exploration-oriented behavior policy without external guidance, which is usually in the form of intrinsic reward, is by increasing the expressiveness of μ to capture information about potentially rewarding actions. Energy-based representations have recently been increasingly chosen as the target form to construct an expressive behavior policy. Since its first introduction by [43] to achieve maximum-entropy reinforcement learning, several additional prior work keeps improving upon this idea. Among them, the most notable ones include Soft Q-Learning (SQL) [14], EntRL [31], and Soft Actor-Critic (SAC) [16]. Collectively, they have achieved competitive results on many benchmark tasks. Formally, the energy-based behavior policy is defined as

$$\mu(a | s) \propto \exp(Q(s, a)), \quad (2)$$

where Q is commonly selected to be the target critic $Q_{\mathcal{R}}^{\pi}$ in prior work [15, 16]. Various efficient samplers have been proposed to approximate the distribution specified in Eq (2). Among them, [14]’s Stein variational gradient descent (SVGD) [22, 38] based sampler is especially worth noting as it has the potential to approximate complex and multi-model behavior policies. Given this, we also choose it to sample the behavior policy in our proposed ADAC.

Additionally, we want to highlight an intriguing property of SVGD that is critical for understanding why we can perform *analogous disentangled exploration* effectively. Intuitively, SVGD transforms a set of particles to match a target distribution. In the context of RL, following Amortized SVGD [7], we use a neural network sampler $f_{\varphi}(s, \xi)$ ($\xi \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$) to approximate Eq (2), which is done

by minimizing the KL divergence between two distributions. According to [7], f_{φ} is updated according to the following gradient:

$$\begin{aligned} \nabla_{\varphi} J_{\mu}(\varphi) \approx \mathbb{E}_{s, \xi \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\sum_{j=1}^K \underbrace{\left[\mathcal{K}(a, a'_j) \nabla_{a'_j} Q(s, a'_j) \right]}_{\text{term 1}} \right. \\ \left. + \beta \cdot \underbrace{\nabla_{a'_j} \mathcal{K}(a, a'_j)}_{\text{term 2}} \right]_{a=f_{\varphi}(s, \xi)} \frac{\partial f_{\varphi}(s, \xi)}{\partial \varphi} \Big/ K, \end{aligned} \quad (3)$$

where \mathcal{K} is a positive definite kernel³, and β is an additional hyper-parameter proposed to make optimality-expressiveness trade-off. The intrinsic connection between Eq (3) and the deterministic policy gradient (i.e. Eq (1)) is introduced in [14] and [7]: the first term of the gradient represents a combination of deterministic policy gradients weighted by the kernel \mathcal{K} , while the second term of the gradient represents an entropy maximization objective.

To aid a better understanding of this relation, we illustrate the distribution approximated by SVGD using different β in a toy example as shown in Figure 1. The dashed line is the approximation target. When β is small, the entropy of the learned distribution is restricted and the overall policy leans towards the highest-probability region. On the other hand, larger β leads to more expressive approximation.

4 METHOD

This section introduces our proposed method *Analogous Disentangled Actor-Critic* (ADAC). We start by providing an overview of it (Section 4.1), which is followed by elaborating the specific choices we make to design our actors and critics (Sections 4.2 and 4.3).

4.1 Algorithm Overview

Figure 2 provides a diagram overview of ADAC, which consists of two pairs of actor-critic $\langle \mu, Q_{\mathcal{R}}^{\pi} \rangle$ and $\langle \pi, Q_{\mathcal{R}}^{\pi} \rangle$ (see the blue and pink box) to achieve *disentanglement*. Same with prior off-policy algorithms (e.g., DDPG), during training ADAC alternates between the two main procedures, namely *sample collection* (dotted green box), where we use μ to interact with the environment to collect training samples, and *model update* (dashed gray box), which consists of two phases: (i) batches of the collected samples are used to update both critics (the pink box); (ii) μ and π (the blue box) are updated according to their respective critic using different objectives. During evaluation, π is used to interact with the environment.

Both steps in the *model update* phase manifest the analogous property of our method. First, although optimized with respect to different objectives, both μ and π are represented by the neural network f , where $\mu(s) := f_{\varphi}(s, \xi) |_{\xi \sim \mathcal{N}(\mathbf{0}, \mathbf{I})}$ and $\pi(s) := f_{\varphi}(s, \xi) |_{\xi = [0, \dots, 0]^T}$.⁴ That is, π is a deterministic policy since its input ξ is fixed, while $\mu(s)$ can be regarded as an action sampler that uses the randomly sampled ξ to generate actions. As we shall demonstrate in Section 4.2, this specific setup effectively restricts the deviation between the two policies (μ and π) (i.e. update bias), which stabilizes the training process and maintains sufficient expressiveness in the behavior policy μ (also see Section 5.1 for an intuitive illustration).

³Formally, in ADAC, we define the kernel as $\mathcal{K}(a, \hat{a}_i) = \frac{1}{\sqrt{2\pi(d/K)}} \exp\left(-\frac{\|a - \hat{a}_i\|^2}{2(d/K)^2}\right)$,

where d is the number of dimensions of the action space.

⁴ f takes two components s and ξ as input, and φ is the parameter set of f .




Figure 2: Block diagram of ADAC, which consists of the sample collection phase (green box with dotted line) and the model update phase (gray box with dashed line). Model (i.e. actor and critic networks) updates are performed sequentially from ① to ④. Each update step’s corresponding line in Algorithm 1 is shown in brackets.

The second exhibit of our method’s *analogous* nature lies on our designed critics Q_R^π and $Q_{R'}^\pi$, which are based on the environment-defined reward \mathcal{R} and the augmented reward $\mathcal{R}' := \mathcal{R} + \mathcal{R}^{\text{in}}$ (\mathcal{R}^{in} is the intrinsic reward) respectively yet are both computed with regard to the target policy π . As a standard approach, Q_R^π approximates the task-defined objective that the algorithm aims to maximize. On the other hand, $Q_{R'}^\pi$ is a behavior critic that can be shown to be both *explorative* and *stable* theoretically (Section 4.3) and empirically (Section 5.3). Note that when not using intrinsic reward, the two critics are degraded to be identical to one another (i.e. $\mathcal{R} = \mathcal{R}'$) and in practice when that happens we only store one of them.

To better appreciate our method, it is not enough to only gain an overview about our actors and critics in isolation. Given this, we then formalize the connections between the actors and the critics as well as the objectives that are optimized during the model update phase (Figure 2). As defined above, π is the exploitation policy that aims to maintain optimality throughout the learning process, which is best optimized using the deterministic policy gradient (Eq (1)), where Q_R^π is used as the referred critic (① in Figure 2). On the other hand, for the sake of expressiveness, the energy-based objective (Eq (2)) is a good fit for μ . To further encourage exploration, we use the behavior critic $Q_{R'}^\pi$ in the objective, which gives $\mu(a | s) \propto \exp(Q_{R'}^\pi(s, a))$ (② in Figure 2). Since both policies share the same network f , the actor optimization process (③ in Figure 2) is done by maximizing

$$J_\pi(\varphi) + J_\mu(\varphi), \quad (4)$$

where the gradients of both terms are defined by Eqs (1) and (3), respectively. In particular, we set $\pi(s) := f_\varphi(s, \xi)|_{\xi=[0,0,\dots,0]^T}$ in Eq (1) and $Q := Q_{R'}^\pi$ in Eq (3). As illustrated in Algorithm 1 (line 5), we update Q_R^π and $Q_{R'}^\pi$ with the target $\mathcal{T}^\pi Q_R^\pi$ and $\mathcal{T}^\pi Q_{R'}^\pi$ on the collected samples using the mean squared error loss, respectively.

In the sample collection phase, μ interacts with the environment and the gathered samples are stored in a replay buffer [24] for later use in the model update phase. Given state s , actions are sampled from μ with a three-step procedure: (i) sample $\xi \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, (ii) plug the sampled ξ in $f_\varphi(s, \xi)$ to get its output \hat{a} , and (iii) regard \hat{a} as the center of kernel $\mathcal{K}(\cdot, \hat{a})^1$ and sample an action a from it.

On the implementation side, ADAC is compatible with any existing off-policy actor-critic model for continuous control: it directly builds upon them by inheriting their actor π (which is also their

Algorithm 1 The model update phase of ADAC. Correspondence with Figure 2 is given after “//”.

- 1: **Input:** A minibatch of samples \mathcal{B} , actor model f_φ (represents the target policy f_φ^π as well as the behavior policy f_φ^μ), critic models Q_R^π and $Q_{R'}^\pi$.
- 2: $\nabla_\varphi f_\varphi^\pi \leftarrow$ the deterministic policy gradient of Q_R^π with respect to π (Eq (1)). // **target policy update**
- 3: $\nabla_\varphi f_\varphi^\mu \leftarrow$ gradient of $Q_{R'}^\pi$ with respect to the behavior policy μ (Eq (3), Section 3.3) // **behavior policy learning**
- 4: Update f with $\nabla_\varphi f_\varphi^\pi$ and $\nabla_\varphi f_\varphi^\mu$ // **policy co-training**
- 5: Update Q_ϕ^π and Q_ψ^μ to minimize the mean squared error on \mathcal{B} with respect to the target $\mathcal{T}^\pi Q_R^\pi$ and $\mathcal{T}^\pi Q_{R'}^\pi$, respectively. // **value update with critic bounding**

target policy) and critic $Q_{R'}^\pi$. To be more specific, ADAC merely adds a new actor μ to interact with the environment and a new critic $Q_{R'}^\pi$ that guides μ ’s updates on top of the base model, along with the constraints/connections enforced between the inherited and the new actor and between the inherent and the new critic (i.e. policy co-training and critic bounding). In other words, modifications made by ADAC would not conflict with the originally proposed improvements on the base model. In our experiments, two base models (i.e. DDPG [21] and TD3 [11]) are adopted.⁵

4.2 Stabilizing Policy Updates by Policy Co-training

Although a behavior policy by Eq (2) is sufficiently expressive to capture potentially rewarding actions, it may still not be helpful for learning a better π : being expressive also means that μ is often significantly different from π , leading to collect samples that can substantially bias π ’s updates (recall the discussion about Equation 1), and in turn rendering the learning process of Q_R^π unstable and vulnerable to catastrophic failure [12, 30, 36, 41]. To be more specific, since the difference between π and an expressive μ is more than some zero-mean random noise, the state marginal distribution (ρ_μ) defined with respect to μ can potentially diverge greatly from that (ρ_π) defined with respect to π . Since ρ_π is not directly accessible, as shown in Eq (1), the gradients of π are approximated using samples from ρ_μ . When the approximated gradients constantly deviate significantly from the true values (i.e. the approximated gradients are biased), the updates to π essentially become inaccurate and hence ineffective. This suggests that a brutal act of disentangling the behavior policy from the target policy alone is not a guarantee of improved training efficiency or final performance.

Therefore, to mitigate the aforementioned problem, we would like to reduce the distance between μ and π , which naturally reduces the KL-divergence between distribution ρ_μ and ρ_π . One straightforward approach to reduce the distance between the two policies is to restrict the randomness of μ , for example by lowering the entropy of the behavior policy μ through a smaller β (Eq (3)). However, this inevitably sacrifices μ ’s expressiveness, which in turn would also harm ADAC’s competitiveness. Alternatively, we propose *policy co-training* to best maintain the expressiveness of μ while also stabilizing it with regards to π , which is motivated by the intrinsic connection between Eqs (1) and (3)

⁵Detailed algorithm tables are available in the longer version of this paper.

(see the 2nd paragraph of Section 3.3). We reiterate here that in a nutshell, both policies are modeled by the same network f and are distinguished only by their different inputs to ξ . During training, f is updated to maximize Eq (4). The method to sample actions from μ is described in the 5th paragraph of Section 4.1.

We further justify the above choice by demonstrating that the imposed restrictions on μ and π only have minor influence on π 's optimality and μ 's expressiveness. To argue for this point, we need to revisit Eq (3) for one more time: π can be viewed as being updated with $\beta = 0$, whereas μ is updated with $\beta > 0$. Intuitively, this makes policy π optimal since its action is not affected by the entropy maximization term (i.e. the second term). μ is still expressive since only when the input random variable ξ is close to the zero vector, it will be significantly restricted by π . In Section 5.1, we will empirically demonstrate policy co-training indeed reduces the distance between μ and π during training, fulfilling its mission.

Additionally, *policy co-training* enforces the underlying relations between π and μ . Specifically, policy co-training forces π to be contained in μ since $[0, 0, \dots, 0]^T$ is the highest-density point of $\mathcal{N}(\mathbf{0}, \mathbf{I})$, and sampling ξ from $\mathcal{N}(\mathbf{0}, \mathbf{I})$ is likely to generate actions close to that from π . This matches the intuition that π and μ should share similarities: actions proposed by π is rewarding (with respect to \mathcal{R}) and thus should be frequently executed by μ .

4.3 Incorporating Intrinsic Reward in Behavior Critic via Critic Bounding

With the help of disentanglement as well as *policy co-training*, we manage to design an expressive behavior policy that not only explores effectively but also helps stabilize π 's learning process. In this subsection, we aim to achieve the same goal – stability and expressiveness – on a different subject, the behavior critic $Q_{\mathcal{R}'}$.

As introduced in Section 4.1, \mathcal{R} is the environment-defined reward function, while \mathcal{R}' consists of an additional exploration-oriented intrinsic reward \mathcal{R}^{in} . As hinted by the notations, ADAC's target critic $Q_{\mathcal{R}}^\pi$ and behavior critic $Q_{\mathcal{R}'}^\pi$ are defined with regard to the same policy but updated differently according to the following

$$Q_{\mathcal{R}}^\pi \leftarrow \mathcal{T}_{\mathcal{R}}^\pi Q_{\mathcal{R}}^\pi; \quad Q_{\mathcal{R}'}^\pi \leftarrow \mathcal{T}_{\mathcal{R}'}^\pi Q_{\mathcal{R}'}^\pi, \quad (5)$$

where updates are performed through minibatches in practice. Note that when no intrinsic reward is used, Eq (5) becomes trivial and the two critics ($Q_{\mathcal{R}}^\pi$ and $Q_{\mathcal{R}'}^\pi$) are identical. Therefore, we only consider the case where intrinsic reward exists in the following discussion.

While it is natural that the target critic is updated using the target policy, it may seem counterintuitive that the behavior critic is also updated using the target policy. Given that μ is updated following the guidance (i.e. through the energy-based objective) of $Q_{\mathcal{R}'}^\pi$, we do so to prevent μ from diverging disastrously from π .

THEOREM 4.1. *Let π be a greedy policy w.r.t. $Q_{\mathcal{R}}^\pi$ and μ be a greedy policy w.r.t. $Q_{\mathcal{R}'}^\pi$. Assume $Q_{\mathcal{R}'}^\pi$ is optimal w.r.t. $\mathcal{T}_{\mathcal{R}'}^\pi$ and $\mathcal{R}'(s, a) \geq \mathcal{R}(s, a)$ ($\forall s, a \in \mathcal{S} \times \mathcal{A}$). We have the following results.*

First, $\mathbb{E}_{\rho_\pi}[\mathcal{T}_{\mathcal{R}}^{\max} Q_{\mathcal{R}}^\pi - Q_{\mathcal{R}}^\pi]$, a proxy of training stability, is lower bounded by

$$\mathbb{E}_{\rho_\mu}[\mathcal{T}_{\mathcal{R}}^{\max} Q_{\mathcal{R}}^\pi - Q_{\mathcal{R}}^\pi] + \mathbb{E}_{\rho_\pi}[\mathcal{R}] - \mathbb{E}_{\rho_\mu}[\mathcal{R}]. \quad (6)$$

Second, $\mathbb{E}_{\rho_\mu}[\mathcal{T}_{\mathcal{R}}^{\max} Q_{\mathcal{R}}^\pi - Q_{\mathcal{R}}^\pi]$, a proxy of training effectiveness, is lower bounded by

$$\mathbb{E}_{\rho_\pi}[\mathcal{T}_{\mathcal{R}}^{\max} Q_{\mathcal{R}}^\pi - Q_{\mathcal{R}}^\pi] + \mathbb{E}_{\rho_\pi}[\mathcal{R} - \mathcal{R}']. \quad (7)$$

We first examine its assumptions. While others are generally satisfiable and are commonly made in the RL literature [26], the assumption on the rewards – $\forall s, a \in \mathcal{S} \times \mathcal{A}, \mathcal{R}'(s, a) \geq \mathcal{R}(s, a)$ – seems restrictive. However, since most intrinsic rewards are strictly greater than zero (e.g., [10, 17]), it can be easily satisfied in practice.

The full proof is deferred to the longer version of this paper. Here, we only focus on the insights conveyed by Theorem 4.1. According to the definition of the Bellman optimality operator (Section 3.1), $\mathbb{E}_{\rho}[\mathcal{T}_{\mathcal{R}}^{\max} Q_{\mathcal{R}}^\pi - Q_{\mathcal{R}}^\pi]$ quantifies the improvement on $Q_{\mathcal{R}}^\pi$ after performing one value iteration step (w.r.t. \mathcal{R}) [3]. Depending on the state-action distribution ρ used to compute expectation, this quantity becomes a proxy of different measures. Specifically, $\mathbb{E}_{\rho_\pi}[\mathcal{T}_{\mathcal{R}}^{\max} Q_{\mathcal{R}}^\pi - Q_{\mathcal{R}}^\pi]$ (where the expectation is calculated w.r.t. ρ_π) represents the expected improvement of the target policy, which is our ultimate learning goal and hence is a proxy of *training stability* given learning is stable if non-decreasing. $\mathbb{E}_{\rho_\mu}[\mathcal{T}_{\mathcal{R}}^{\max} Q_{\mathcal{R}}^\pi - Q_{\mathcal{R}}^\pi]$ (where the expectation is with regard to ρ_μ) measures the training *effectiveness* in the sense that the better μ performs (i.e. $\mathbb{E}_{\rho_\mu}[Q_{\mathcal{R}}^\pi]$), the higher the quality of collected samples, given μ is used to interact with the environment.

Since both proxies measure the near-future improvement of $Q_{\mathcal{R}}^\pi$, higher values are preferable. In particular, if of both the lower bounds are guaranteed to be positive, ADAC's performance will be monotonously increasing during training even with intrinsic reward. Note that the lower bound of both proxies are inter-defined on each other through their first terms, we are not able to comment on the exact value of either alone. Still, we can share several encouraging observations, which are related to the remaining parts of both bounds. First, the remaining part of Eq (6) (i.e. $\mathbb{E}_{\rho_\pi}[\mathcal{R}] - \mathbb{E}_{\rho_\mu}[\mathcal{R}]$) is always non-negative since π is optimized to maximize the cumulative reward of \mathcal{R} while μ is not. This suggests that even if $\mathbb{E}_{\rho_\mu}[\mathcal{T}_{\mathcal{R}}^{\max} Q_{\mathcal{R}}^\pi - Q_{\mathcal{R}}^\pi]$ is negative, Eq (6) might still be positive, which guarantees training stability. Next, we demonstrate that the remaining term of Eq (7) (i.e. $\mathbb{E}_{\rho_\pi}[\mathcal{R} - \mathcal{R}']$) is very likely to be high during ADAC's training process. Promised by the policy co-training approach (Section 4.2), when following μ we would frequently visit high-probability state-action pairs in ρ_π . Since most intrinsic reward functions are designed to assign small rewards to states that are frequently visited, $\mathbb{E}_{\rho_\pi}[\mathcal{R}' - \mathcal{R}] = \mathbb{E}_{\rho_\pi}[\mathcal{R}^{in}]$ (the negation of the remaining term of Eq (7)) would be small. Finally, bringing the two bounds together, we can show that when π and μ are jointly optimized with the presence of intrinsic rewards, the performance (w.r.t. the environment-defined reward \mathcal{R}) of both policies is very unlikely to drop catastrophically, thanks to the remaining terms of both bounds, which are either positive or very likely to be high.

5 EXPERIMENTS

In this section, we take gradual steps to analyze and illustrate our proposed method ADAC. Specifically, We first investigate the behavior of our analogous disentangled behavior policy μ (Section 5.1). Next, we perform an empirical evaluation of ADAC without intrinsic rewards on 14 standard continuous-control benchmarks (Section 5.2). Finally, encouraged by its promising performance and to further justify the critic bounding method, we examine ADAC with intrinsic rewards in 4 sparse-reward and hence exploration-heavy environments (Section 5.3). Throughout this paper, we highlight




Figure 3: Learning curves of ADAC (with base model DDPG) and baselines on the modified CartPole environment. In addition, ADAC’s target (red dots) and behavior policies (solid blue curves) at different timesteps are plotted below the learning curves.

two benefits from the analogous disentangled nature of ADAC: (i) avoiding unnecessary trade-offs between current optimality and exploration (i.e. a more expressive and effective behavior policy); (ii) natural compatibility with intrinsic rewards without altering environment-defined optimality. In this context, the first two subsections are devoted to demonstrating the first benefit and the last subsection is dedicated for the second.

5.1 Analysis of Analogous Disentangled Behavior Policy

Since we are largely motivated by the potential luxury of designing an expressive exploration strategy offered by the disentangled nature of our framework, it is natural we are first interested in investigating how well our behavior policy lives up to this expectation. Yet as discussed in Section 4.2, in order to aid stable policy updates, we specifically put some restraints on our behavior policy, deliberately making it analogous of the target policy, which means our behavior policy may not be as expressive as otherwise. Given this, we start with investigating whether our behavior policy is still expressive enough, which is measured by its *coverage* (i.e. does it explore a wide enough action/policy space outside the current target policy). To further examine the influence of our added constraints, we study the policy network’s *stability* (i.e. does the *policy co-training* lowers the bias between two policies and stabilize π ’s learning process). Finally, we focus on the effectiveness of our behavior policy by measuring the overall *performance* of ADAC (i.e. does ADAC’s exploration strategy efficiently lead the target policy to iteratively converge to a more desirable local optimum).

Setup For the sake of easy illustration, we choose a straightforward environment, namely CartPole [5], as our demonstration bed. The goal in this environment is to balance the pole that is attached to a cart by applying left/right force. For the compatibility with continuous control and a better modeling of real-world implications, we modify CartPole’s original discrete action space and add effort penalty to the rewards as specified in Table 1. To demonstrate the advantages of our behavior policy, we choose DDPG with two commonly-used existing exploration strategies as the main baselines, i.e., Gaussian noise (σ_g) and Ornstein-Uhlenbeck process noise (σ_o) [37], both with two variance levels 0.2 and 1.0.

For the virtue of fair comparison, we only present DDPG-based ADAC here (or simply ADAC later in this subsection). To further demonstrate the benefits from disentanglement, we choose SAC as another baseline. As discussed earlier in related works, SAC similarly also utilizes energy-based policies, yet opposite to our approach its exploration is embedded into its target policy.

Empirical Insights To minimize distraction, our discussion starts with closely examining ADAC’s behavior and target policy alone. First, see the cells at the bottom of Figure 3, which are snapshots of the behavior and target policy at different training stages. As suggested by the wide bell shape of the solid blue curves (μ) at the first cell, our behavior policy acts curiously when ignorant about the environment, extensively exploring all possible actions including those that are far away from the target policy (represented by the red dots). Yet having such a broad coverage alone is still not sufficient to overcome the beginning trap of getting stuck in the deceiving local-optimum of constantly exerting $a \in [-0.5, 0.5]$. As suggested by the bimodal shape of the solid blue curve (μ) in the second cell, after acquiring a preliminary understanding of the environment the agent starts to form preference for some actions when exploring. Almost at the same time, the target policy no longer stays close to 0.0 (represented by the intersection of the two axes), suggesting that the behavior policy is effective in leading the target policy towards a more desirable place. This can be further corroborated by what is suggested from the third and fourth cell. In the late stage, besides being able to balance the pole, our agent even manages to learn exerting actions with small absolute value from time to time to avoid the effort penalty.

Other than its expressiveness, stability critically influences ADAC’s overall performance, which is by design controlled by the proposed *policy co-training* approach. To examine its effect, we perform an ablation study about it. To be more specific, we compare ADAC with ADAC without *policy co-training*.⁶ The effect of critic bounding is measured by the *bias* between π and μ , which is shown in the middle of Figure 3. We can see that ADAC has much lower bias than its variant without policy co-training. Additionally, policy co-training does not affect the expressiveness of μ , which is suggested by the behavior policies rendered below in Figure 3.

Finally we move our attention to the learning curves in Figure 3: ADAC exceeds baselines in both learning efficiency (i.e. being the first to consistently accumulate positive rewards) and final performance. Unlike our behavior policy, exploration through random noise is unguided, resulting in either wasted exploration on unpromising regions or insufficient exploration on rewarding areas. This largely explains the noticeable performance gap between DDPG with random noise and ADAC. On the other side, SAC bears an expressive policy similar to our behavior policy. However, suffering from no separate behavior policy, to aid exploration, SAC has to consistently take sub-optimal actions into account, adversely affecting its policy improvement process. In other words, different from ADAC, SAC cannot fully exploits its learned knowledge of the environment (i.e. its value functions) to construct its target policy, leading to a performance inferior to ADAC’s.

⁶When policy co-training is not used, we make two copies of ADAC’s original policy network (which encodes both π and μ) to represent π and μ , respectively. Gradients of the policies are applied only to their corresponding networks.

Table 1: Specifications of our action and reward designs for the modified CartPole task. The original task consists of two discrete actions *left* and *right*, each pushing the cart towards its corresponding direction. We converted them into a single-dimension continuous action.

Action ($a \in [-1, 1]$)	Reward ($r \in \mathbb{R}$)
$a = \begin{cases} \text{left} & a < -0.5 \\ p(\text{left}) = p(\text{right}) = 0.5 & a \in [0.5, 0.5] \\ \text{right} & a > 0.5 \end{cases}$	$r = -0.1 a - 0.05a^2 + \begin{cases} -1.0 & \text{episode ended} \\ 0.1 & \text{otherwise} \end{cases}$

Table 2: Continuous-control performance in 14 benchmark environments. Average episode return (\pm standard deviation) over 20 trials are reported. Bold indicates the best average episode return. \dagger indicates the better performance between ADAC(TD3) and its base model TD3. Similarly, $*$ indicates the better performance between ADAC (DDPG) and its base model DDPG. In all three cases, values that are statistically insignificantly different (>0.05 in t-test) from the respective should-be indicated ones are denoted as well.

Environment	ADAC (TD3)	ADAC (DDPG)	TD3	DDPG	SAC	PPO
RoboschoolAnt	2219 \pm 373	838.1 \pm 97.1	2903 $\dagger \pm$ 666	450.0 \pm 27.9	2726 \pm 652	1280 \pm 71
RoboschoolHopper	2299 $\dagger \pm$ 333	766.5 \pm 10	2302 $\dagger \pm$ 537	543.8 \pm 307	2089 \pm 657	1229 \pm 345
RoboschoolHalfCheetah	1578 $\dagger \pm$ 166	1711 \pm 95	607.2 \pm 246.2	441.6 \pm 120.4	807.0 \pm 252.6	1225 \pm 184.2
RoboschoolAtlasForwardWalk	234.6 $\dagger \pm$ 55.7	186.7 \pm 37.9	190.6 \pm 50.1	52.63 \pm 26.2	126.0 \pm 47.1	107.6 \pm 29.4
RoboschoolWalker2d	1769 $\dagger \pm$ 452	1564 \pm 651	995.1 \pm 146.3	208.7 \pm 137.1	1021 \pm 263	578.9 \pm 231.3
Ant	3353 \pm 847	1226 \pm 18	4034 $\dagger \pm$ 517	370.5 \pm 223	4291 \pm 1498	1401 \pm 168
Hopper	3598 $\dagger \pm$ 374	374.5 \pm 36.5	2845 \pm 609	38.93 \pm 0.88	3307 \pm 825	1555 \pm 458
HalfCheetah	9392 \pm 199	2238 \pm 40	10526 $\dagger \pm$ 2367	1009 \pm 49	11541 \pm 2989	881.7 \pm 10.1
Walker2d	5122 $\dagger \pm$ 1314	1291 \pm 42	4630 $\dagger \pm$ 778	186.2 \pm 33.3	4067 \pm 1211	1146 \pm 368
InvertedPendulum	1000 $\dagger \pm$ 0	1000 \pm 0	1000 $\dagger \pm$ 0	1000 \pm 0	1000 \pm 0	98.90 \pm 2.08
InvertedDoublePendulum	9359 $\dagger \pm$ 0.17	9334 \pm 1.39	7665 \pm 566	27.20 \pm 2.61	9353 \pm 2896	98.90 \pm 5.88
BipedalWalker	309.8 $\dagger \pm$ 15.6	-52.77 \pm 1.94	288.4 $\dagger \pm$ 51.25	-123.90 \pm 11.17	307.2 \pm 57.92	266.9 \pm 28.52
BipedalWalkerHardcore	- 10.76 $\dagger \pm$ 27.70	-98.52 \pm 3.21	-57.97 \pm 21.08	-50.05 \pm 10.27	-127.4 \pm 45.2	-105.3 \pm 22.2
LunarLanderContinuous	290.0 $\dagger \pm$ 50.9	85.67 \pm 23.42	289.7 $\dagger \pm$ 54.1	-65.89 \pm 96.48	283.3 \pm 69.29	59.32 \pm 68.44

5.2 Comparison with the State of the Art

Though well-suited for illustration, CartPole alone is not challenging and generalized enough to fully manifest ADAC’s competitiveness. In this subsection, we present that ADAC can achieve state-of-the-art performance in standard benchmarks.

Setup To demonstrate the generality of our method, we construct a 14-task testbed suite composed of qualitatively diverse continuous-control environments from the OpenAI Gym toolkit [5]. On top of the two baselines adopted earlier (i.e. DDPG and SAC), we further include TD3 [11], which improves upon DDPG by addressing some of its function approximation errors, PPO [32], which is regarded as one of the most stable and efficient on-policy policy gradient algorithm, and GEP-PG [6], which combines Goal Exploration Process [28] with policy gradient to perform curious exploration as well as stable learning. Though not exhaustive, this baseline suite still embodies many of the latest advancements and can be indeed deemed as the existing state-of-the-art. However, we compare with GEP-PG only in tasks adopted in their original experiments. Specifically, since the GEP part of the algorithm needs hand-crafted exploration goals, we are not able to run their model on new experiments since it is nontrivial to generalize their experiments in other tasks. To best reproduce the rest’s performance, we use their original open-source implementations if released; otherwise, we build our own versions after the most-starred third-party implementations in GitHub. Furthermore, to prevent over-claiming the state-of-the-art, we fine-tune their hyper-parameters around the values reported in the respective literature, but only coarsely tune the hyper-parameters introduced by ADAC.⁷ All experiments are run for 1 million time-steps, or until reaching performance convergence, whichever happens earlier.

Empirical Insights Table 2 corroborates that ADAC’s competitiveness over existing methods stem from its disentangled nature. More importantly, these results reveal two desirable properties of

ADAC’s full compatibility with existing off-policy methods. First, ADAC consistently outperforms the method it is based on. As indicated by the $*$ symbols, compared to its base model, DDPG-based ADAC achieves statistically better or comparable performance on more than 93%(13/14) of the benchmarks and obtains identical performance on one of the remaining two. Though not as remarkable as DDPG-based ADAC, TD-based ADAC also manages to achieve statically better or comparable performance over its base model on more than 78%(11/14) of the tasks; see the \dagger symbols. Second, ADAC retains the benefits of improvements developed by the base model themselves. This is best illustrated by TD3-based ADAC’s performance superiority over DDPG-based ADAC.

We would like to specially call readers’ attention to our comparison of ADAC with SAC since they both use energy-based behavior policy. This comparison also reveals the benefit brought by the disentangled structure and the analogous actors and critics. ADAC (TD3) achieves better average performance over SAC on 71%(10/14) of the benchmarks, indicating the effectiveness of our proposed analogous disentangled structure.

Despite that GDP-PG [6] also uses the disentanglement idea, we do not compare ADAC with it across the whole 14-benchmark test suite and hence GEP-PG is not included in Table 2. This decision is made largely due to the fact that the Goal Exploration Process (GEP) in GEP-PG requires manually defining a goal space to explore, which is task dependent and can critically influence the algorithm performance. Given this, we can only compare with it on the two experiments that GEP-PG has run, of which only one overlaps with our task suit, namely HalfCheetah. In HalfCheetah, GEP-PG achieves 6118 cumulative reward, while ADAC (TD3) achieves 9392, showing superiority over GEP-PG. Furthermore as also acknowledged in its paper, GEP-PG lags behind SAC in performance, which

⁷For a fair comparison, we do not use intrinsic reward throughout this section, since most baseline approaches are not able to naturally incorporate them during learning.




Figure 4: Illustration of how intrinsic reward contaminates the environment-defined optimality in PendulumSparse. Fooled into collecting more intrinsic rewards rather than environment rewards (see the learning curves on the left), the agent constantly alternates between spinning the pendulum and barely moving it (see the snapshots of the target policy on the right), making no real progress.

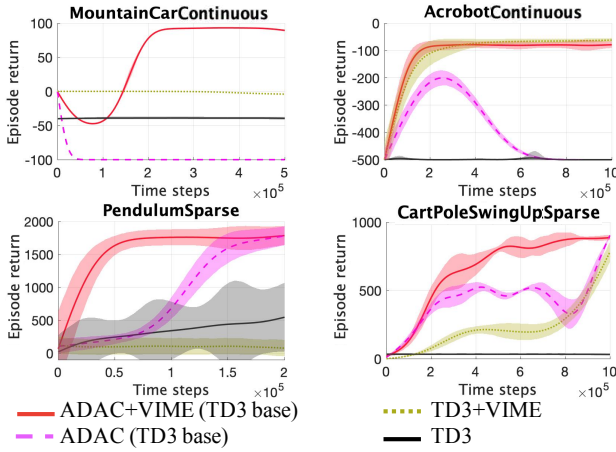


Figure 5: Learning curves for four sparse-reward tasks. Lines denote the average over 20 trials and the shaded areas represent the range of one standard deviation.

suggests that naively disentangling the behavior policy from the target policy does not guarantee competitive performance. Rather, to design effective disentangled actor-critic, we should also pay attention to how to best restrict some components.

When considering all reported methods together, TD3-based ADAC obtains the most number of the state-of-the-art results; as indicated in bold, it is the best performer (or statistically comparable with the best) on more than 71%(10/14) of the benchmarks.

5.3 Evaluation in Sparse-Reward Environments

Encouraged by the promising results observed on the benchmarks, in this subsection we evaluate ADAC under more challenging environments, in which rewards are barely provided. This set of experiments aim to test ADAC’s exploration capacity under extreme settings. Furthermore, we also see them fit as demonstration beds to present ADAC’s natural compatibility with intrinsic methods. In this regard, we are particularly interested in investigating whether the disentangled nature of ADAC helps mitigate intrinsic rewards’ undesirable bias effect on the environment-defined optimality.

Setup To the surprise of many, sparse-reward environments turn out to be relatively unpopular in commonly-used RL toolkits. Besides including the classic MountainCarContinuous and Acrobot

(after converting its action space to be continuous), to construct a decently sized testing suite, we further hand-craft new tasks, namely PendulumSparse and CartPoleSwingUpSparse by sparsifying the rewards in the existing environments. It is achieved mainly through suppressing the original rewards until reaching some predefined threshold.⁸ Due to their dependency on environment-provided rewards as feedback signals, most model-free RL algorithms suffer significant performance degradation in these sparse-reward tasks. In this situation, resorting to intrinsic methods (IM) for additional signals has been widely considered as the go-to solution. Among a wide variety of IM methods, we adopt Variational Information Maximization Exploration (VIME) [17] as our internal reward generator for its consistent good performance on a wide variety of exploration-challenging tasks. Considering TD3-based ADAC’s superiority over DDPG-based ADAC, we only combine VIME into TD3 and TD3-based ADAC. Note when paired with ADAC, intrinsic rewards are only visible to the behavior policy.

Empirical Insights Among the four environments, PendulumSparse has the most vulnerable environment-defined optimality. The goal here is to swing the inverted pendulum up so it stays upright. As suggested by Figure 4, not knowing how to distinguish between intrinsic and environment rewards, VIME-augmented TD3 is completely fooled into chasing after the intrinsic rewards. In other words, the VIME-augmented TD3’s understanding of what is optimal is completely off from the true environment-defined optimality. Note as demonstrated in left-bottom part of Figure 5, VIME-augmented TD3’s performance even trails behind TD3’s, which is an indisputable evidence that the bias introduced by IM can be detrimental and should be addressed whenever possible. In contrast, thanks to its disentangled nature, VIME-augmented ADAC only perceives intrinsic rewards in its behavior policy, which means its target policy always remains optimal with regards to our current knowledge about environment rewards. Because of this, VIME-augmented manages to consistently solve this exploration-challenging task. ADAC’s natural compatibility with VIME is further corroborated by the results in the remaining 3 tasks. As suggested by the complete Figure 5, VIME-augmented ADAC consistently surpasses all reported alternatives by a large margin in terms of both convergence speed and final performance.

6 CONCLUSION

We present Analogous Disentangled Actor-Critic (ADAC), an off-policy reinforcement learning framework that explicitly disentangles the behavior and target policy. Compared to prior work, to stabilize model updates, we restrain our behavior policy and its corresponding critic to be analogous of their target counterparts. Thanks to its disentangled and analogous nature, environment-reward-only ADAC achieves the state-of-the-art results in 10 out of 14 continuous control benchmarks. Moreover, ADAC is naturally compatible with intrinsic rewards, outperforming alternatives in exploration-challenging tasks.

Acknowledgements This work is partially supported by NSF grants #IIS-1943641, #IIS-1633857, #CCF-1837129, DARPA XAI grant #N66001-17-2-4032, UCLA Samueli Fellowship, and gifts from Intel and Facebook Research.

⁸More details about the sparse-reward environments are reported in the longer version.

REFERENCES

- [1] Kai Arulkumaran, Marc Peter Deisenroth, Miles Brundage, and Anil Anthony Bharath. 2017. Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine* 34, 6 (2017), 26–38.
- [2] Marc Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi Munos. 2016. Unifying count-based exploration and intrinsic motivation. In *Advances in Neural Information Processing Systems*. 1471–1479.
- [3] Richard Bellman. 1966. Dynamic programming. *Science* 153, 3731 (1966), 34–37.
- [4] Lucas Beyer, Damien Vincent, Olivier Teboul, Sylvain Gelly, Matthieu Geist, and Olivier Pietquin. 2019. MULEX: Disentangling Exploitation from Exploration in Deep RL. *arXiv preprint arXiv:1907.00868* (2019).
- [5] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. OpenAI Gym. (2016). *arXiv:arXiv:1606.01540*
- [6] Cédric Colas, Olivier Sigaud, and Pierre-Yves Oudeyer. 2018. Gep-pg: Decoupling exploration and exploitation in deep reinforcement learning algorithms. *arXiv preprint arXiv:1802.05054* (2018).
- [7] Yihao Feng, Dilin Wang, and Qiang Liu. 2017. Learning to draw samples with amortized stein variational gradient descent. *arXiv preprint arXiv:1707.06626* (2017).
- [8] Sébastien Forestier, Yoan Mollard, and Pierre-Yves Oudeyer. 2017. Intrinsically motivated goal exploration processes with automatic curriculum learning. *arXiv preprint arXiv:1708.02190* (2017).
- [9] Meire Fortunato, Mohammad Gheshlaghi Azar, Bilal Piot, Jacob Menick, Ian Osband, Alex Graves, Vlad Mnih, Remi Munos, Demis Hassabis, Olivier Pietquin, et al. 2017. Noisy networks for exploration. *International Conference on Learning Representations (ICLR)* (2017).
- [10] Justin Fu, John Co-Reyes, and Sergey Levine. 2017. Ex2: Exploration with exemplar models for deep reinforcement learning. In *Advances in Neural Information Processing Systems*. 2577–2587.
- [11] Scott Fujimoto, Herke Hoof, and David Meger. 2018. Addressing Function Approximation Error in Actor-Critic Methods. In *International Conference on Machine Learning*.
- [12] Scott Fujimoto, David Meger, and Doina Precup. 2018. Off-policy deep reinforcement learning without exploration. *arXiv preprint arXiv:1812.02900* (2018).
- [13] Shixiang Gu, Timothy Lillicrap, Zoubin Ghahramani, Richard E Turner, and Sergey Levine. 2016. Q-prop: Sample-efficient policy gradient with an off-policy critic. *arXiv preprint arXiv:1611.02247* (2016).
- [14] Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. 2017. Reinforcement learning with deep energy-based policies. In *Proceedings of the International Conference on Machine Learning-Volume 70*.
- [15] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning*.
- [16] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. 2018. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905* (2018).
- [17] Rein Houthoof, Xi Chen, Yan Duan, John Schulman, Filip De Turck, and Pieter Abbeel. 2016. Vime: Variational information maximizing exploration. In *Advances in Neural Information Processing Systems*. 1109–1117.
- [18] Tejas D Kulkarni, Karthik Narasimhan, Ardavan Saeedi, and Josh Tenenbaum. 2016. Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation. In *Advances in neural information processing systems*. 3675–3683.
- [19] Joel Z. Leibo, Vinicius Zambaldi, Marc Lanctot, Janusz Marecki, and Thore Graepel. 2017. Multi-agent Reinforcement Learning in Sequential Social Dilemmas. In *Proc. of AAMAS*.
- [20] Yitao Liang, Marlos C. Machado, Erik Talvitie, and Michael Bowling. 2016. State of the Art Control of Atari Games Using Shallow Reinforcement Learning. In *Proc. of AAMAS*.
- [21] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. 2015. Continuous control with deep reinforcement learning. In *Proceedings of the International Conference on Learning Representations*.
- [22] Qiang Liu and Dilin Wang. 2016. Stein variational gradient descent: A general purpose bayesian inference algorithm. In *Advances In Neural Information Processing Systems*. 2378–2386.
- [23] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. 2016. Asynchronous methods for deep reinforcement learning. In *Proceedings of the International conference on machine learning*. 1928–1937.
- [24] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602* (2013).
- [25] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharmashan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. 2015. Human-level Control through Deep Reinforcement Learning. *Nature* 518, 7540 (26 02 2015), 529–533.
- [26] Rémi Munos. 2007. Performance bounds in L_p -norm for approximate value iteration. *SIAM journal on control and optimization* 46, 2 (2007), 541–561.
- [27] Ofir Nachum, Mohammad Norouzi, Kelvin Xu, and Dale Schuurmans. 2018. Trust-PCL: An Off-Policy Trust Region Method for Continuous Control. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=HyrCWeWcb>
- [28] Alexandre Péré, Sébastien Forestier, Olivier Sigaud, and Pierre-Yves Oudeyer. 2018. Unsupervised learning of goal spaces for intrinsically motivated goal exploration. *arXiv preprint arXiv:1803.00781* (2018).
- [29] Matthias Plappert, Rein Houthoof, Prafulla Dhariwal, Szymon Sidor, Richard Y Chen, Xi Chen, Tamim Asfour, Pieter Abbeel, and Marcin Andrychowicz. 2018. Parameter space noise for exploration. In *Proceedings of the International Conference on Learning Representations*.
- [30] Matthew Schlegel, Wesley Chung, Daniel Graves, Jian Qian, and Martha White. 2019. Importance Resampling for Off-policy Prediction. *arXiv preprint arXiv:1906.04328* (2019).
- [31] John Schulman, Xi Chen, and Pieter Abbeel. 2017. Equivalence between policy gradients and soft q-learning. *arXiv preprint arXiv:1704.06440* (2017).
- [32] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- [33] David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. 2014. Deterministic Policy Gradient Algorithms. In *International Conference on International Conference on Machine Learning*.
- [34] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. 2017. Mastering the game of Go without human knowledge. *Nature* 550 (18 10 2017), 354 EP –.
- [35] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- [36] Richard S Sutton, Csaba Szepesvári, and Hamid Reza Maei. 2008. A convergent $O(n)$ algorithm for off-policy temporal-difference learning with linear function approximation. *Advances in neural information processing systems* 21, 21 (2008), 1609–1616.
- [37] G. E. Uhlenbeck and L. S. Ornstein. 1930. On the Theory of the Brownian Motion. *Phys. Rev.* 36 (1930), 823–841. Issue 5.
- [38] Dilin Wang, Zhe Zeng, and Qiang Liu. 2018. Stein Variational Message Passing for Continuous Graphical Models. In *ICML*. 5206–5214. <http://proceedings.mlr.press/v80/wang18l.html>
- [39] Yue Wang and Fumin Zhang. 2017. *Trends in Control and Decision-Making for Human-Robot Collaboration Systems* (1st ed.). Springer Publishing Company, Incorporated.
- [40] Christopher J. C. H. Watkins and Peter Dayan. 1992. Technical Note: Q-Learning. *Machine Learning* 8, 3-4 (May 1992).
- [41] Tengyang Xie, Yifei Ma, and Yu-Xiang Wang. 2019. Optimal Off-Policy Evaluation for Reinforcement Learning with Marginalized Importance Sampling. *arXiv preprint arXiv:1906.03393* (2019).
- [42] Zhi-Xiong Xu, Xi-Liang Chen, Lei Cao, and Chen-Xi Li. 2017. A study of count-based exploration and bonus for reinforcement learning. In *Proceedings of the IEEE International Conference on Cloud Computing and Big Data Analysis (ICCCBDA)*. IEEE, 425–429.
- [43] Brian D Ziebart. 2010. *Modeling purposeful adaptive behavior with the principle of maximum causal entropy*. Ph.D. Dissertation. figshare.