

## Offline Reinforcement Learning as Anti-exploration

Shideh Rezaeifar,<sup>1</sup> Robert Dadashi,<sup>2</sup> Nino Vieillard,<sup>2,3</sup> Léonard Hussenot,<sup>2,4</sup> Olivier Bachem,<sup>2</sup>  
Olivier Pietquin,<sup>2</sup> Matthieu Geist<sup>2</sup>

<sup>1</sup> University of Geneva

<sup>2</sup> Google Research, Brain Team

<sup>3</sup> Université de Lorraine, CNRS, Inria, IECL, F-54000 Nancy, France

<sup>4</sup> Université de Lille, CNRS, Inria, UMR 9189 CRISAL

shideh.rezaeifar@unige.ch, dadashi@google.com, vieillard@google.com, hussenot@google.com, bachem@google.com, pietquin@google.com, mfggeist@google.com

### Abstract

Offline Reinforcement Learning (RL) aims at learning an optimal control from a fixed dataset, without interactions with the system. An agent in this setting should avoid selecting actions whose consequences cannot be predicted from the data. This is the converse of exploration in RL, which favors such actions. We thus take inspiration from the literature on bonus-based exploration to design a new offline RL agent. The core idea is to subtract a prediction-based exploration bonus from the reward, instead of adding it for exploration. This allows the policy to stay close to the support of the dataset, and practically extends some previous pessimism-based offline RL methods to a deep learning setting with arbitrary bonuses. We also connect this approach to a more common regularization of the learned policy towards the data. Instantiated with a bonus based on the prediction error of a variational autoencoder, we show that our simple agent is competitive with the state of the art on a set of continuous control locomotion and manipulation tasks.

### Introduction

Deep Reinforcement Learning (RL) has achieved remarkable success in a variety of tasks including robotics (Kober, Bagnell, and Peters 2013; Thrun 1995; Benbrahim and Franklin 1997; Bagnell and Schneider 2001; Endo et al. 2008), recommendation systems (Rojanavasu, Srinil, and Pinngern 2005; Afsar, Crump, and Far 2021), and games (Silver et al. 2018). Deep RL algorithms generally assume that an agent repeatedly interacts with an environment and uses the gathered data to improve its policy: they are said to be *online*. Because actual interactions with the environment are necessary to online RL algorithms, they do not comply with most of real-world applications constraints. Indeed, allowing an agent to collect new data may be infeasible such as in healthcare (Murphy et al. 2001), autonomous driving (Sallab et al. 2017; Grigorescu et al. 2020), or education (Mandel et al. 2014). As an alternative, *offline* RL (Levine et al. 2020) is a practical paradigm where an agent is trained using a fixed dataset of trajectories, without any interactions with the environment during learning. The ability to learn from a fixed dataset is the crucial step towards scalable and generalizable data-driven learning methods.

In principle, off-policy algorithms (Ernst, Geurts, and Wehenkel 2005; Lagoudakis and Parr 2003; Mnih et al. 2013; Lillicrap et al. 2016; Haarnoja et al. 2018) could be used to learn from a fixed dataset. However, in practice, they perform poorly without any feedback from the environment. This issue persists even when the off-policy data comes from effective expert policies, which in principle should address any exploration issue (Kumar et al. 2019). The main challenge comes from the sensitivity to the data distribution. The distribution mismatch between the behavior policy and the learned policy leads to an extrapolation error of the value function, which can become overly optimistic in areas of state-action space outside the support of the dataset. The extrapolation error accumulates along the episode and results in unstable learning and divergence (Fujimoto, Hoof, and Meger 2018; Kumar et al. 2019, 2020; Levine et al. 2020; Peng et al. 2019; Dadashi et al. 2021; Yu et al. 2020).

This work introduces a new approach to Offline RL, inspired by exploration. This may seem counter-intuitive. Indeed, in online RL, an exploring agent will try to visit state-action pairs it has never experienced before, hence drifting from the distribution in the dataset. This is exactly what an Offline RL agent should avoid, so we frame it as an *anti-exploration* problem. We focus on bonus-based exploration (Brafman and Tennenholtz 2002; Burda et al. 2019; Pathak et al. 2017; Burda et al. 2018; Bellemare et al. 2016b). The underlying principle, in these methods, is to add a bonus to the reward function, this bonus being higher for novel/surprising state action-pairs (Barto, Mirolli, and Baldassarre 2013). The core idea of the proposed approach is to perform *anti-exploration* by *subtracting* the bonus from the reward, instead of adding it, effectively preventing the agent from selecting actions of which the effects are unknown. Concurrently, Rashidinejad et al. (2021) proposed a theoretical framework for subtracting an exploration bonus from the reward that was applied only to the tabular setting. Similarly, Jin, Yang, and Wang (2021) recently proposed a theoretical pessimistic variant of the value iteration based on uncertainty quantifier. However, they did not provide any practical deep RL algorithm with empirical results. Under minimal assumptions, we relate this approach to the more common offline RL idea of penalizing the learned policy for deviating from the policy that supposedly generated the data. We also propose a specific instantiation of this general idea, using TD3 (Fuji-

moto, Hoof, and Meger 2018) as the learning agent, and for which the bonus is the reconstruction error of a variational autoencoder (VAE) trained on offline state-action pairs. We evaluate the agent on the hand manipulation and locomotion tasks of the D4RL benchmark (Fu et al. 2020), and show that it is competitive with the state of the art.

## Preliminaries

A Markov decision process (MDP) is defined by a tuple  $\mathcal{M} := (\mathcal{S}, \mathcal{A}, r, P, \gamma)$ , with  $\mathcal{S}$  the state space,  $\mathcal{A}$  the action space,  $P \in \Delta_{\mathcal{S}^{\times \mathcal{A}}}$  the Markovian transition kernel ( $\Delta_{\mathcal{S}}$  is the set of distributions over  $\mathcal{S}$ ),  $r \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$  the reward function and  $\gamma \in (0, 1)$  the discount factor. A policy  $\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}$  is a mapping from states to distribution over actions (a deterministic policy being a special case). The general objective of RL is to find a policy that maximizes the expectation of the return  $G_t = \sum_{i=0}^{\infty} \gamma^i r(s_{t+i}, a_{t+i})$ . Given a policy  $\pi$ , the action value function, or Q-function, is defined as the expected return when following the policy  $\pi$  after taking action  $a$  in state  $s$ ,  $Q^\pi(s, a) = \mathbb{E}_\pi[G_t | s_t = s, a_t = a]$ .

For  $u_1, u_2 \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ , define the dot product  $\langle u_1, u_2 \rangle = \sum_a u_1(\cdot, a) u_2(\cdot, a) \in \mathbb{R}^{\mathcal{S}}$ , and for  $v \in \mathbb{R}^{\mathcal{S}}$  write  $Pv = \sum_{s'} P(s' | \cdot, \cdot) v(s') \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ . Whith these notations, define the Bellman operator as  $\mathcal{B}^\pi Q = r + \gamma P \langle \pi, Q \rangle$ . It is a contraction (Puterman 2014) and  $Q^\pi$  is its unique fixed point. An optimal policy  $\pi^*$  satisfies  $\pi^* \in \operatorname{argmax}_\pi Q^\pi$ . A policy  $\pi$  is said to be a greedy policy with respect to  $Q$  if  $\pi \in \operatorname{argmax}_{\pi'} \langle \pi', Q \rangle$  (notice that  $\langle \pi, Q \rangle(s) = E_{a \sim \pi(\cdot | s)}[Q(s, a)]$ ). An algorithm that allows computing an optimal policy is Value Iteration (VI):

$$\begin{cases} \pi_{k+1} \in \operatorname{argmax}_{\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}} \langle \pi, Q_k \rangle \\ Q_{k+1} = r + \gamma P \langle \pi_{k+1}, Q_k \rangle \end{cases}$$

The first step is the greedy step. When considering only deterministic policies (slightly abusing notations), it simplifies to  $\pi_{k+1}(s) = \operatorname{argmax}_{a \in \mathcal{A}} Q_k(s, a)$ . The second step is the evaluation step, again for deterministic policies, it simplifies to  $Q_{k+1}(s, a) = r(s, a) + \gamma \mathbb{E}_{s' | s, a}[\max_{a'} Q_k(s', a')]$ .

VI may be used to derive many existing deep RL algorithms. First, consider the discrete action case, assume that a dataset of collected transitions  $\mathcal{D} = \{(s, a, s', r)\}$  is available. The seminal DQN (Mnih et al. 2013) parameterizes the Q-value with a deep neural network  $Q_\omega$ , takes  $Q_k$  as being a copy of a previous network  $Q_{\bar{\omega}}$  ( $\bar{\omega}$  being a copy of the parameters), uses the fact that the greedy policy can be exactly computed, and considers a loss which is the squared difference of both sides of the evaluation step,  $L_{\text{dqn}}(\omega) = \hat{\mathbb{E}}_{\mathcal{D}}[(r + \gamma \max_{a'} Q_{\bar{\omega}}(s', a') - Q_\omega(s, a))^2]$ , with  $\hat{\mathbb{E}}$  an empirical expectation.

With continuous actions, the greedy policy can no longer be computed exactly, which requires introducing a policy network  $\pi_\theta$ . TD3 (Fujimoto, Hoof, and Meger 2018), a state-of-the-art actor-critic algorithm, can be derived from the same VI viewpoint. It adopts a fixed-variance Gaussian parameterization for the policy,  $\pi_\theta \sim \mathcal{N}(\mu_\theta, \sigma I)$  with  $\mu_\theta \in \mathcal{A}^{\mathcal{S}}$ ,  $\sigma$  the standard deviation of the exploration noise and  $I$  the identity matrix. The greedy step can be approximated by the actor loss  $J_{\text{td3, actor}}(\theta) = \hat{\mathbb{E}}_{s \in \mathcal{D}}[\hat{\mathbb{E}}_{a \sim \pi_\theta(\cdot, s)}[Q_{\bar{\omega}}(s, a)]]$ . This can

be made more convenient by using the reparameterization trick,  $J_{\text{td3, actor}}(\theta) = \hat{\mathbb{E}}_{s \in \mathcal{D}}[\hat{\mathbb{E}}_{\epsilon \sim \mathcal{N}(0, \sigma I)}[Q_{\bar{\omega}}(s, \mu_\theta(s) + \epsilon)]]$ . The critic loss is similar to the DQN one,  $J_{\text{td3, critic}}(\omega) = \hat{\mathbb{E}}_{\mathcal{D}}[\hat{E}_{a' \sim \pi_\theta(\cdot | s')}[(r + \gamma Q_{\bar{\omega}}(s', a') - Q_\omega(s, a))^2]$ .

## Framing Offline RL as Anti-Exploration

**What makes offline reinforcement learning hard?** Conventional on-policy RL algorithms operate in a setting where an agent repeatedly interacts with the environment, gathers new data and uses the data to update its policy. The term off-policy denotes an algorithm that can use the data collected by other policies whilst still being able to interact with the environment. Offline RL relies solely on a previously collected dataset without further interaction with the environment. This setting can leverage the vast amount of previously collected datasets, *e.g.* human demonstrations or hand-engineered exploration strategies. Offline RL is challenging because the collected dataset does not cover the entire state-action space. The out-of-distribution (OOD) actions can cause an extrapolation error in the value function approximations. As an example, consider the regression target of the DQN loss,  $y = r + \gamma \max_{a' \in \mathcal{A}} Q_{\bar{\omega}}(s', a')$ . The estimate of the value function  $Q_{\bar{\omega}}(s', a')$  for state-action pairs that are not in the dataset can be erroneously high due to extrapolation errors. As a consequence, the maximum  $\max_{a' \in \mathcal{A}} Q_{\bar{\omega}}(s', a')$  may be reached for a state-action couple  $(s', a')$  that has never been observed in the data. Using this maximal value as part of the target for estimating  $Q_\omega(s, a)$  will result in being over-optimistic about  $(s, a)$ . This estimation error accumulates with time as the policy selects actions that maximize the value function. Thus, many methods constrain the learned policy to stay within the support of the dataset. These methods differ in how the deviation is measured and how the constraint is enforced. For example, it could be achieved by modifying the DQN target to  $r + \gamma \max_{a' | (s', a') \in \mathcal{D}} Q_\omega(s', a')$ , which is the underlying idea of Kumar et al. (2019); Ghasemipour, Schuurmans, and Gu (2020), among others.

**Exploration in Reinforcement Learning.** Exploration is critical in RL and failing at handling it properly can prevent agents to identify high reward regions, or even to gather any reward at all if it is sparse. There are many approaches to exploration, as well as many challenges, such as the well-known exploration-exploitation dilemma. In this paper, we focus on bonus-based exploration and then adapt it to offline RL. The core idea is to define or learn a bonus function  $b \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ , which is small for known state-action pairs and high for unknown ones. This bonus is *added* to the reward function which will intuitively drive the learned policy to follow trajectories of unknown state-action pairs. Embedded within a generic VI approach, this can be written as

$$\begin{cases} \pi_{k+1} \in \operatorname{argmax}_{\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}} \langle \pi, Q_k \rangle \\ Q_{k+1} = r + b + \gamma P \langle \pi_{k+1}, Q_k \rangle \end{cases}$$

The bonus-based strategies can be roughly categorized into count-based or prediction-based methods. First, in count-based methods, the novelty of a state is measured by the number of visits, and a bonus is assigned accordingly. For example, it can be  $b(s, a) \propto 1/\sqrt{n(s, a)}$ , with  $n(s, a)$

the number of times the state-action couple has been encountered. The bonus guides the agent’s behavior to prefer rarely visited states over common ones. When the state-action space is too large, counting is not a viable option, but the frequency of visits can be approximated by using a density model (Bellemare et al. 2016b; Ostrovski et al. 2017) or mapping state-actions to hash codes (Tang et al. 2017). Second, in prediction-based methods, the novelty is related to the agent’s knowledge about the environment. The agent’s familiarity with the environment can be estimated through a prediction model (Achiam and Sastry 2017; Schmidhuber 1991; Pathak et al. 2017; Burda et al. 2018). For example, a forward prediction model captures an agent’s ability to predict the outcome of its own actions. High prediction error indicates that the agent is less familiar with that state-action pair and vice versa. However, this prediction error is entangled with the agent’s performance in the specific task. Random network distillation (RND) was introduced as an alternative method where the prediction task is random and independent of the main one (Burda et al. 2019).

**Proposed anti-exploration approach.** Offline RL algorithms learn from a fixed dataset without any interaction with the environment. State-action pairs outside of the dataset are therefore never actually experienced and can receive erroneously optimistic value estimates because of extrapolation errors that are not corrected by an environment feedback. The overestimation of value functions can be encouraged in online reinforcement learning, to a certain extent, as it incentivizes agents to explore and learn by trial and error (Gulcehre et al. 2021; Schmidhuber 1991). Moreover, in an online setting, if the agent wrongly assigns a high value to a given action, this action will be chosen, the real return will be experimented, and the value of the action will be corrected through bootstrapping. In this sense, online RL is self-correcting. On the contrary, in offline RL, the converse of exploration is required to keep the state-actions close to the support of the dataset (no self-correction is possible given the absence of interaction). Hence, a natural idea consists in defining an *anti-exploration* bonus to penalize OOD state-action pairs. This bonus will encourage the policy to act similarly to existing transitions of the offline trajectories.

A naive approach to anti-exploration for offline RL would thus consist in *subtracting* the bonus from the reward, instead of adding it. Embedded within our general VI scheme, it is

$$\begin{cases} \pi_{k+1} = \operatorname{argmax}_{\pi} \langle \pi, Q_k \rangle \\ Q_{k+1} = r - b + \gamma P \langle \pi_{k+1}, Q_k \rangle \end{cases} .$$

Intuitively, this should prevent the RL agent from choosing actions with high bonus, i.e., unknown actions that are not or not enough present in the dataset. However, this would not be effective at all. Indeed, in offline RL, we would only use state-action pairs in the dataset, for which the bonus is supposedly low. Hence this would not avoid bootstrapping values of actions with unknown consequences, which is our primary goal. As an example, consider again the bootstrapped target of DQN, with the additional bonus,  $r - b(s, a) + \gamma \max_{a' \in \mathcal{A}} Q_{\bar{\omega}}(s', a')$ . The state action-couple  $(s, a)$  necessarily comes from the dataset, so the bonus is supposedly low, for example zero for a well trained prediction-

based bonus. Thus, the bonus is here essentially useless. What would make more sense would be to penalize the bootstrapped value with the bonus, for example considering the target  $r + \gamma \max_{a' \in \mathcal{A}} (Q_{\bar{\omega}}(s', a') - b(s', a'))$ . This way, we penalize bootstrapping values for unknown state-action pairs  $(s', a')$ . It happens conveniently that both approaches are equivalent, from a dynamic programming viewpoint (that is, when Q-values and policies are computed exactly for all possible state-action pairs; the equivalence obviously does not hold in an approximate offline setting). Indeed, we have

$$\begin{aligned} & \begin{cases} \pi_{k+1} = \operatorname{argmax}_{\pi} \langle \pi, Q_k \rangle \\ Q_{k+1} = r - b + \gamma P \langle \pi_{k+1}, Q_k \rangle \end{cases} \\ \Leftrightarrow & \begin{cases} \pi_{k+1} = \operatorname{argmax}_{\pi} \langle \pi, Q'_k - b \rangle \\ Q'_{k+1} = r + \gamma P \langle \pi_{k+1}, Q'_k - b \rangle \end{cases} , \end{aligned} \quad (1)$$

with  $Q'_{k+1} = Q_{k+1} + b$ . Even though the  $Q$ -values are not the same, both algorithms provide the same sequence of policies, provided that  $Q'_0 = Q_0 + b$ . In fact, independently from the initial  $Q$ -values, both algorithms will converge to the policy maximizing  $\mathbb{E}_{\pi} [\sum_t \gamma^t (r(s_t, a_t) - b(s_t, a_t))]$ . This comes from the fact that this is a specific instance of a regularized MDP (Geist, Scherrer, and Pietquin 2019), with a linear regularizer  $\Omega(\pi) = \langle \pi, b \rangle$ . This is a better basis for an offline agent, as the bonus directly affects the  $Q$ -function. This was illustrated above with the example of the DQN target as a specific instance of this VI scheme. However, the idea holds more generally, for example within an actor-critic scheme, as we will exemplify later by providing an agent based on TD3 for continuous actions.

## A Link to Regularization

Many recent papers in offline RL focus on regularizing the learned policy to stay close to the training dataset of offline trajectories.

It usually amounts to a penalization based on a divergence between the learned policy and the behavior policy (the action distribution conditioned on states underlying the dataset). These methods share the same principle, but they differ notably in how the deviation of policies is measured. Here, we draw a link between anti-exploration and behavior-regularized offline RL.

Let  $b \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$  be any exploration bonus. We will just assume that the bonus is lower for state-action pairs in the dataset than for those outside. For example, if  $b$  was trained with a one-class classifier, it could be  $b(s, a) \approx 0$  if  $(s, a)$  in the support,  $b(s, a) \approx 1$  elsewhere. We will discuss a different approach later, based on the reconstruction error of a conditional variational autoencoder.

We can use this bonus to model a distribution of actions conditioned on states, assigning high probabilities for low bonuses, and low probabilities for high bonuses, the goal being to model the data distribution (which is a hard problem in general). This can be conveniently done with a softmax distribution. Let  $\beta > 0$  be a scale parameter and  $\tau > 0$  a temperature, we define the policy  $\pi_b = \operatorname{softmax}(-\frac{\beta}{\tau} b(s, \cdot))$ . Now, we can use this policy to regularize a VI scheme. Define the Kullback-Leibler (KL) divergence between two policies

as  $\text{KL}(\pi_1||\pi_2) = \langle \pi_1, \ln \pi_1 - \ln \pi_2 \rangle \in \mathbb{R}^S$ . Define also the entropy of a policy as  $\mathcal{H}(\pi) = -\langle \pi, \ln \pi \rangle \in \mathbb{R}^S$ . Consider the following KL-regularized VI scheme (Geist, Scherrer, and Pietquin 2019; Vieillard et al. 2020):

$$\begin{cases} \pi_{k+1} = \operatorname{argmax}_{\pi} (\langle \pi, Q_k \rangle - \tau \text{KL}(\pi||\pi_b)) \\ Q_{k+1} = r + \gamma P (\langle \pi_{k+1}, Q_k \rangle - \tau \text{KL}(\pi_{k+1}||\pi_b)). \end{cases} \quad (2)$$

Consider for example the quantity optimized within the greedy step, we have that  $\langle \pi, Q_k \rangle - \tau \text{KL}(\pi||\pi_b) = \langle \pi, Q_k + \tau \ln \operatorname{softmax}(-\frac{\beta b}{\tau}) \rangle - \tau \langle \pi, \ln \pi \rangle = \langle \pi, Q_k - \Delta_b \rangle + \mathcal{H}(\pi)$ , with  $\Delta_b(s, a) = \beta b(s, a) + \tau \ln \sum_{a'} \exp(-\beta b(s, a')/\tau)$ . The same derivation can be done for the evaluation step:

$$(2) \Leftrightarrow \begin{cases} \pi_{k+1} = \operatorname{argmax}_{\pi} (\langle \pi, Q_k - \Delta_b \rangle + \tau \mathcal{H}(\pi)) \\ Q_{k+1} = r + \gamma P (\langle \pi_{k+1}, Q_k - \Delta_b \rangle + \tau \mathcal{H}(\pi_{k+1})). \end{cases}$$

This can be seen as an entropy-regularized variation of the scheme we proposed Eq. (1). Now, taking the limit as the temperature goes to zero, we have  $\lim_{\tau \rightarrow 0} \Delta_b(s, a) = \beta b(s, a) + \max_{a'} (-\beta b(s, a')) = \beta(b(s, a) - \min_{a'} b(s, a'))$ .

Assuming moreover that the bonus is a prediction-based bonus, well-trained, meaning that  $\min_{a'} b(s, a') = 0$  for any state  $s$  in the dataset (as there is an action associated to it), we can rewrite the limiting case as Eq. (1):

$$\begin{cases} \pi_{k+1} = \operatorname{argmax}_{\pi} \langle \pi, Q_k - b \rangle \\ Q_{k+1} = r + \gamma P \langle \pi_{k+1}, Q_k - b \rangle \end{cases}$$

Thus, we can see the proposed anti-exploration VI scheme as a limiting case, when the temperature goes to zero, of a KL-regularized scheme that regularize the learned policy towards a behavioral policy constructed from the bonus, assigning higher probabilities to lower bonuses. This derivation is reminiscent to advantage learning (Bellemare et al. 2016a) being a limiting case of KL-regularized VI (Vieillard et al. 2020) in online RL (Vieillard, Pietquin, and Geist 2020).

In a tabular setting, choosing the bonus as  $b(s, a) = -\frac{\tau}{\beta} \ln n(s, a)$ , we would obtain a frequency-based approximation of the behavioral policy. Yet, this does not scale easily, and we propose a more practical approach.

## A Practical Approach

In principle, any (VI-based) RL agent could be combined with any exploration bonus for providing an anti-exploration agent for offline RL. For example, in a discrete action setting, one could combine DQN, described briefly before, with RND (Burda et al. 2019). The RND exploration bonus is defined as the prediction error of encoded features as shown in Figure 1: a prediction network is trained to predict the output of a fixed random neural network denoted as target network, ideally leading to higher prediction errors when facing unfamiliar states. Here, we will focus on a continuous action setting. To do so, we consider the TD3 actor-critic as the RL agent, briefly described before. For the exploration bonus, RND is a natural choice as it is a well-performing bonus for the Atari benchmark and has been extended to continuous state-action support estimation in the context of Imitation Learning with RED (Wang et al. 2019). However, in our experiments, RND performed poorly.

Note that RND was introduced for discrete action spaces and image-based observations, for which random CNNs capture meaningful features. In the continuous state-action setting, random MLPs do not capture environment specific features (in RED for instance, the random networks are tuned for each environment, and are used together with a BC regularization term). We show in Section that RND is not discriminative enough to distinguish state-action pairs in the dataset from others. Therefore, we introduce a bonus based on the reconstruction error of a variational autoencoder.

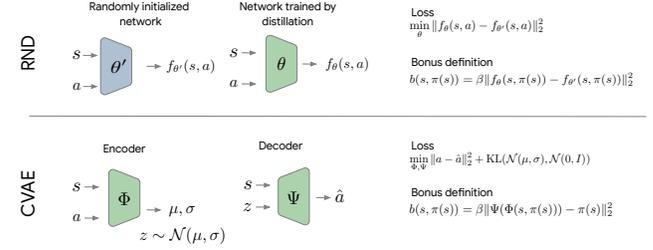


Figure 1: Illustration of RND and CVAE networks, losses and inferred anti-exploration bonuses.

**TD3.** We described briefly TD3 before, from a VI viewpoint. It indeed comes with additional particularities, that are important for good empirical performance. The noise (the standard deviation of the Gaussian policy) is not necessarily the same in the actor critic losses, a twin-critic is considered for reducing the overestimation bias when bootstrapping, and policy updates are delayed. We refer to Fujimoto, Hoof, and Meger (2018) for more details. We use the classic TD3 update, except for the additional bonus term:

$$J_{\text{td3, actor}}, b(\theta) = \hat{\mathbb{E}}_{s \in \mathcal{D}} [Q_{\bar{\omega}}(s, \mu_{\theta}(s)) - b(s, \mu_{\theta}(s))], \quad (3)$$

$$J_{\text{td3, critic}}, b(\omega) = \hat{\mathbb{E}}_{\mathcal{D}} [\hat{E}_{a' = \mu_{\theta}(s'), \epsilon \sim \mathcal{N}(0, \sigma I)} (r + \gamma(Q_{\bar{\omega}}(s', a' + \epsilon) - b(s', a' + \epsilon)) - Q_{\omega}(s, a))^2]. \quad (4)$$

**CVAE.** The bonus we use for anti-exploration is based on a Conditional Variational Autoencoder (CVAE) (Sohn, Lee, and Yan 2015). The Variational Autoencoder (VAE) was first introduced by Kingma and Welling (2014). The model consists of two networks, the encoder  $\Phi$  and the decoder  $\Psi$ . The input data  $x$  is encoded to a latent representation  $z$  and then the samples  $\hat{x}$  are generated by the decoder from the latent space. Let us consider a dataset,  $X = \{x^1, \dots, x^N\}$ , consisting of  $N$  i.i.d. samples. We assume that the data were generated from low dimensional latent variables  $z$ . VAE performs density estimation on  $P(x, z)$  to maximize the likelihood of the observed training data  $x$ :  $\log P(x) = \sum_{i=1}^N \log P(x^i)$ . Since this marginal likelihood is difficult to work with directly for non-trivial models, instead a parametric inference model  $\Psi(z|x)$  is used to optimize the variational lower bound on the marginal log-likelihood:  $\mathcal{L}_{\Phi, \Psi} = E_{P(z|x)} [\log \Phi(x|z)] - \text{KL}(\Psi(z|x)||\Phi(z))$ . A VAE optimizes the lower bound by reparameterizing  $\Psi(z|x)$  (Kingma and Welling 2014; Rezende, Mohamed, and Wierstra 2014). The

first term of  $\mathcal{L}$  corresponds to the reconstruction error and the second term regularizes the distribution parameterized by the encoder  $\Phi(x|z)$  to minimize the KL divergence from a chosen prior distribution, usually an isotropic, centered Gaussian. In our problem formulation, given a dataset  $\mathcal{D}$ , we use a conditional variational autoencoder to reconstruct actions conditioned on the states. Hence,  $\mathcal{L}_{\Phi, \Psi}$  is equal to

$$E_{P(z|s,a)}[\log \Phi(a|s,z)] - \text{KL}(\Psi(z|s,a) \parallel \Phi(z|s)). \quad (5)$$

**Summary.** Our specific instantiation of the idea of anti-exploration for offline RL works as follows. Given a dataset of interactions, we train a CVAE to predict actions conditioned on states (Alg. 1). For any given state-action pair, the bonus is the scaled prediction error of the CVAE,

$$b(s,a) = \beta \|a - \Psi(\Phi(s,a))\|_2^2, \quad (6)$$

with  $\beta$  the scale parameter. This bonus modifies the TD3 losses, run over the fixed dataset (Alg. 2).

---

Algorithm 1: CVAE training.

---

- 1: Init. CVAE networks  $\Phi$  and  $\Psi$
  - 2: **for** step  $i = 0$  to  $N$  **do**
  - 3:   Sample a minibatch of  $k$  state-action pairs from  $\mathcal{D}$
  - 4:   Train  $\Phi$  and  $\Psi$  using  $\mathcal{L}_{\Phi, \Psi}$ , see Eq. (5)
  - 5: **end for**
- 

---

Algorithm 2: Modified TD3 training.

---

- 1: Init. policy  $\pi_\theta$ , Q-network  $Q_\omega$  and target network  $Q_{\bar{\omega}}$
  - 2: **for** step  $i = 0$  to  $N$  **do**
  - 3:   Sample a minibatch of  $k$  transitions from  $\mathcal{D}$
  - 4:   For each transition, compute the bonus, Eq. (6)
  - 5:   Update critic: gradient step on  $J_{\text{id3, critic}, b}$ , Eq. (4)
  - 6:   Update actor: gradient step on  $J_{\text{id3, actor}, b}$ , Eq. (3)
  - 7:   Update target network  $Q_{\bar{\omega}} := Q_\omega$
  - 8: **end for**
- 

## Experiments

After describing the experimental setup and datasets, we first evaluate the discriminative power of the CVAE-based anti-exploration bonus in identifying OOD state-action pairs, comparing it to the one of the more natural (in an exploration context, at least) RND (Sec. Anti-exploration bonus). Then, we compare the proposed approach to prior offline RL methods on a range of hand manipulation and locomotion tasks with multiple data collection strategies (Fu et al. 2020).

**Experimental setup.** We focus on locomotion and manipulation tasks from the D4RL dataset (Fu et al. 2020). Along with different tasks, multiple data collection strategies are also considered for testing the agent’s performance in complex environments. First, for the locomotion tasks, the goal is to maximize the traveled distance. For these tasks, the datasets are: `random`, `medium-replay`, `medium` and lastly `medium-expert`. `Random` consists of transitions collected by a random policy. `Medium-replay` contains the first million transitions collected by a SAC

agent (Haarnoja et al. 2018) trained from scratch on the environment. `Medium` has transitions collected by a policy with sub-optimal performance. Lastly, `medium-expert` is build up from transitions collected by a near optimal policy next to transitions collected by a sub-optimal policy. Second, the hand manipulation tasks require controlling a 24-DoF simulated hand in different tasks of hammering a nail, opening a door, spinning a pen, and relocating a ball (Rajeswaran et al. 2017). These tasks are considerably more complicated than the gym locomotion tasks with higher dimensionality. The following datasets were collected on hand manipulation tasks: `human`, `cloned`, and `expert`. The `human` dataset is collected by a human operator. `Clone` contains transitions collected by a policy trained with behavioral cloning interacting with the environment next to initial demonstrations. Lastly, `expert` is build up from transitions collected by a fine-tuned RL policy interacting in the environment.

**Anti-exploration bonus.** We analyze the quality of the learned bonus in detail for different algorithms. In particular, we are interested in the capability of the anti-exploration bonus to separate state-action pairs in the dataset from others. However, even though it is straightforward to have positive examples (these are those in the dataset), it is much more difficult to define what negative examples are (otherwise, the bonus could be simply trained using a binary classifier). In these experiments, for a state-action pair  $(s,a) \in \mathcal{D}$ , we define an OOD action  $\tilde{a}$  in three different ways. First, we consider actions uniformly drawn from the action space,  $\tilde{a} \sim \mathcal{U}(\mathcal{A})$ . Second, we consider actions from the dataset perturbed with Gaussian noise,  $\tilde{a} = a + \gamma \mathcal{N}(0, I)$ . Third, we consider randomly shuffled actions (for a set of state-action pairs, we shuffle the actions and not the states, which forms new pairs considered as negative examples).

We investigate the discriminative power of the bonus in distinguishing OOD state-action pairs for two different models, namely RND (as it would be a natural choice in an exploration context, at least in a discrete action setting) and CVAE (the proposed approach for learning the bonus).

As for the case of RND (Burda et al. 2019), state-action pairs are passed to the target network  $f$  and prediction network  $f'$ . The prediction network is trained to predict the encoded feature from the target network given the same state-action pair, i.e., minimize the expected MSE between the encoded features. All the implementation details are provided in the Appendix. The bonus is defined as the prediction error of encoded features:  $b(s,a) := \beta \|f(s,a) - f'(s,a)\|_2^2$ . In the CVAE model, a state-action pair  $(s,a)$  is concatenated and encoded to a latent representation  $z$ . This latent representation  $z$  next to state  $s$  is passed to the decoder to reconstruct action  $a$ . Both encoder and decoder consist of two hidden layers of size 750 with the latent size set to 12. We provide further details on the implementation in the Appendix. The bonus is defined as the reconstruction error of the actions:  $b(s,a) := \beta \|a - \Psi(\Phi(s,a))\|_2^2$ . Notice that the CVAE bonus requires sampling the decoder, and thus is random. Yet, empirically, the associated standard deviation is much smaller than the difference between in-distribution and OOD samples (relative scale of roughly 1%).

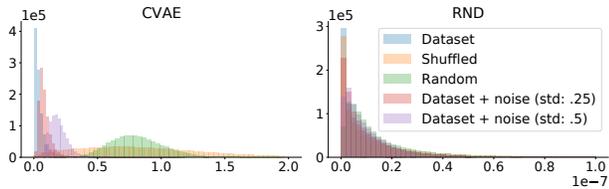


Figure 2: Visualization of the histogram of the reconstruction error for walker2d-medium. The reconstruction error is computed for the original dataset state-action pairs (blue) and different perturbations of the actions: randomly permuted actions over the dataset (orange), random actions (green), original actions to which is added Gaussian noise with different standard deviations (red, purple).

Histograms of the bonus for OOD state-action pairs are compared with those in the dataset and visualized in Figure 2. For these experiments, the state was fixed, and the bonus was derived for different kind of OOD actions. The results are shown for CVAE and RND models in the left and right figures, respectively. RND model performs poorly in identifying OOD actions, and there is not much difference between the bonus for actions in the dataset and randomly generated actions. This might be surprising as Wang et al. (2019) empirically demonstrated that RED reward is sufficiently informative for continuous control tasks in the context of imitation learning. However, estimating the support of an expert policy is different from a non-optimal behavioral policy. Moreover, in RED, environment-specific network tuning is required to achieve good performance, and the reward function comes with a BC regularization term. As we expected, for the CVAE model, the bonus is mostly higher for shuffled, random, and noisy actions in contrast to the actions in the dataset. Moreover, the bonus gets higher as we increase the variance of the added noise. Hence, the CVAE model is discriminative enough to separate state-action pairs in the dataset from others. In fact, in the novelty detection task where the goal is to identify outliers, autoencoders were shown to be beneficial (Abati et al. 2019; Pidhorskyi et al. 2018; Xu et al. 2015; Zhou and Paffenroth 2017).

**Performance on D4RL datasets.** Now that we have assessed the efficiency of CVAE in discriminating OOD state-action pairs from the ones in the dataset, we combine it with TD3 and assess the performance of the resulting offline RL agent on the D4RL datasets depicted above. We compare to the model-free approaches BEAR (Kumar et al. 2019), BRAC (Wu, Tucker, and Nachum 2019), AWR (Peng et al. 2019), BCQ (Fujimoto, Meger, and Precup 2019) and CQL (Kumar et al. 2020), the later providing SOTA results.

The architecture of the TD3 actor and critic consists of a network with two hidden layers of size 256, the first layer has a tanh activation and the second layer has an elu activation. The actor outputs actions with a tanh activation, which is scaled by the action boundaries of each environment. Except from the activation functions, we use the default parameters of TD3 from the authors implementation, and run  $10^6$  gradi-

ent steps using the Adam optimizer, with batches of size 256. Other details are provided in the Appendix.

In online RL, it is quite standard to scale or clip rewards, which can be hard as the reward range may not be known a priori. In an offline setting, all rewards that will be used are in the dataset, so their range is known and it is straightforward to normalize them. Therefore, we scale the reward of each environment such that it is in  $[0, 1]$ . This is important for having a scale parameter  $\beta$  for the bonus not too much dependent from the task. Indeed, we have shown that our anti-exploration idea ideally optimizes for  $\mathbb{E}_\pi[\sum_{t \geq 0} \gamma^t (r(s_t, a_t) - b(s_t, a_t))]$ . Thus, the scale of the bonus should be consistent with the scale of the reward, which is easier with normalized rewards.

TD3 allows for different levels of noises in the actor and critic losses, which deviates a bit from the VI viewpoint. We adopt a similar approach regarding the scale of the bonus, as it provides slightly better results. We allow for different scales  $\beta_a$  and  $\beta_c$  for the bonus in respectively the actor loss and critic loss. We execute a hyperparameter search over the bonus weights over  $\beta_a, \beta_c \in \{0.1, 0.5, 1, 5, 10\}$ . For all locomotion tasks, we select the pair providing the best result (so a single pair for all tasks, not one per task), the same for manipulation tasks. For locomotion  $\beta_a = 5$  and  $\beta_c = 1$ , and for manipulation tasks  $\beta_a = \beta_c = 10$  were chosen.

We show the performance of the proposed approach, TD3-CVAE, on the D4RL datasets in Table 1. We report the average and standard deviation of the results over 10 seeds, each being evaluated on 10 episodes. On average, it is competitive with CQL and outperforms others in locomotion tasks. On the hand manipulation tasks, CVAE outperforms all other methods. Notice that all considered baselines are model-free. Better results can be achieved by model-based approaches (Yu et al. 2021, 2020), at least on the locomotion tasks. Notice that the model-free or model-based aspect is orthogonal to our core contribution, the idea of anti-exploration could be easily combined to a model-based approach. We let further investigations of this aspect for future works.

## Related Work

Here, we briefly discuss prior work in offline RL and how our proposed approach differs. As discussed previously, offline RL suffers from extrapolation error cause by distribution mismatch. Recently, there has been some progress in offline RL to tackle this issue. They can be roughly categorized into policy regularization-based and uncertainty-based methods. The former constrains the learned policy to be as close as possible to the behavioral policy in the dataset. The constraint can be implicit or explicit and ensures that the value function approximator will not encounter out-of-distribution (OOD) state-actions. The difference lies in how the measure of closeness is defined and enforced. Some of the most common measures of closeness are the KL-divergence, maximum mean discrepancy (MMD) distance, or Wasserstein distance (Wu, Tucker, and Nachum 2019). In AWR (Peng et al. 2019), CRR (Wang et al. 2020) or AWAC (Nair et al. 2020), the constraint is introduced implicitly by incorporating a policy update that keeps it close to the behavioral policy. Furthermore, the constraints can be enforced directly on the actor update or the value func-

Algorithm	BC	BEAR	BRACp	BRACv	AWR	BCQ	CQL	TD3-CVAE
halfcheetah-random	2.1	25.1	24.1	31.2	2.5	2.2	<b>35.4</b>	28.6 ± 2.0
walker2d-random	1.6	<b>7.3</b>	-0.2	1.9	1.5	4.9	7.0	5.5 ± 8.0
hopper-random	9.8	11.4	11.0	<b>12.2</b>	10.2	10.6	10.8	11.7 ± 0.2
halfcheetah-medium	36.1	41.7	43.8	<b>46.3</b>	37.4	40.7	44.4	43.2 ± 0.4
walker2d-medium	6.6	59.1	77.5	<b>81.1</b>	17.4	53.1	79.2	68.2 ± 18.7
hopper-medium	29.0	52.1	32.7	31.1	35.9	54.5	<b>58.0</b>	<u>55.9 ± 11.4</u>
halfcheetah-med-rep	38.4	38.6	45.4	<b>47.7</b>	40.3	38.2	46.2	45.3 ± 0.4
walker2d-med-rep	11.3	19.2	-0.3	0.9	15.5	15.0	<b>26.7</b>	15.4 ± 7.8
hopper-med-rep	11.8	33.7	0.6	0.6	28.4	33.1	<b>48.6</b>	46.7 ± 17.9
halfcheetah-med-exp	35.8	53.4	44.2	41.9	52.7	64.7	62.4	<b>86.1 ± 9.7</b>
walker2d-med-exp	6.4	40.1	76.9	81.6	53.8	57.5	<b>111.0</b>	84.9 ± 20.9
hopper-med-exp	<b>111.9</b>	96.3	1.9	0.8	27.1	110.9	98.7	<u>111.6 ± 2.3</u>
Mean performance	25.0	39.8	29.8	31.4	26.8	40.4	<b>52.3</b>	50.3 ± 8.3
pen-human	34.4	-1.0	8.1	0.6	12.3	<b>68.9</b>	37.5	<u>59.2 ± 14.3</u>
hammer-human	1.5	0.3	0.3	0.2	1.2	0.5	<b>4.4</b>	0.2 ± 0.0
door-human	0.5	-0.3	-0.3	-0.3	0.4	-0.0	<b>9.9</b>	0.0 ± 0.0
relocate-human	0.0	-0.3	-0.3	-0.3	-0.0	-0.1	<b>0.2</b>	-0.0 ± 0.0
pen-cloned	<b>56.9</b>	26.5	1.6	-2.5	28.0	44.0	39.2	<u>45.4 ± 25.5</u>
hammer-cloned	0.8	0.3	0.3	0.3	0.4	0.4	<b>2.1</b>	0.3 ± 0.1
door-cloned	-0.1	-0.1	-0.1	-0.1	0.0	0.0	<b>0.4</b>	0.0 ± 0.1
relocate-cloned	<b>-0.1</b>	-0.3	-0.3	-0.3	-0.2	-0.3	-0.1	-0.2 ± 0.0
pen-expert	85.1	105.9	-3.5	-3.0	111.0	<b>114.9</b>	107.0	<u>112.3 ± 21.9</u>
hammer-expert	125.6	127.3	0.3	0.3	39.0	107.2	86.7	<b>128.9 ± 1.5</b>
door-expert	34.9	<b>103.4</b>	-0.3	-0.3	102.9	99.0	101.5	59.4 ± 34.7
relocate-expert	101.3	98.6	-0.3	-0.4	91.5	41.6	95.0	<b>106.4 ± 5.0</b>
Mean performance	36.7	38.3	0.4	-0.4	32.2	39.6	<b>40.3</b>	<b>42.6 ± 8.6</b>

Table 1: Baselines results are reported from Fu et al. (2020), which do not incorporate std of performances, as based on 3 seeds. Following Henderson et al. (2018), we use 10 seeds and evaluate on 10 episodes per seed and report average and std of performance. We bold best average performance and underline the performance when within one std of the best performance.

tion update. In BRAC, Wu, Tucker, and Nachum (2019) investigated different divergences and choices of value penalty or policy regularization. In BEAR, Kumar et al. (2019) argue that restricting the support of the learned policy to the support of the behavior distribution is sufficient, allowing more flexibility and a wider range of actions to the algorithm. An alternative approach to alleviate the effects of OOD state-actions is to make the value function approximators robust to those state-actions. The aim is to make the target value function for OOD state-actions more conservative. This is the underlying principle of CQL (Kumar et al. 2020).

The uncertainty-based approaches for offline RL can build upon robust MDPs (Petrik and Subramanian 2014; Ghavamzadeh, Petrik, and Chow 2016; Larocche, Trichelair, and Des Combes 2019) or upon the concept of pessimism (Yu et al. 2020; Jin, Yang, and Wang 2021; Buckman, Gelada, and Bellemare 2021). The idea there is to adapt the concept of optimism in the face of uncertainty, widely used in exploration, to pessimism for learning from a fixed dataset of transitions. As such, our contribution is conceptually very close to this line of work. Especially, Jin, Yang, and Wang (2021) recently introduced pessimistic VI, which from an abstract viewpoint is very close to what we propose in Eq. (1). They offer theoretical guarantees for linear MDPs, in a finite horizon setting and for a specific bonus, but without empirical study. What we propose, compared to this line of work, is a

practical generalization of this general idea to a deep learning setting (though without theoretical guarantees), allowing the use in principle of any possible exploration bonus. We also introduced a practical and simple bonus built upon CVAE, more efficient than RND, and more practical than uncertainty quantifiers considered in previous works (Yu et al. 2020; Jin, Yang, and Wang 2021; Buckman, Gelada, and Bellemare 2021). We also draw connections between these two general families of offline RL algorithms, regularization-based and uncertainty-based.

## Conclusion

We proposed an intuitive and straightforward approach to offline RL. We constraint the policy to take state-action pairs within the dataset to avoid extrapolation errors. To do so, the core idea is to *subtract* a prediction-based *exploration bonus* from the reward (instead of adding it for exploration). We theoretically showed the connection of the proposed method with the regularization-based approaches. Instantiating this idea with a CVAE-based bonus and the TD3 agent (which is a possibility among many others), we reached competitive performance on D4RL datasets. Even though the performance doesn’t define a new state-of-the-art, we believe that our approach is quite versatile, simple and well principled. As such, an interesting research direction would be to combine it with other orthogonal improvements to offline RL.

## References

- Abati, D.; Porrello, A.; Calderara, S.; and Cucchiara, R. 2019. Latent space autoregression for novelty detection. In *CVPR*, 481–490.
- Achiam, J.; and Sastry, S. 2017. Surprise-based intrinsic motivation for deep reinforcement learning. In *Deep RL Workshop, NeurIPS*.
- Afsar, M. M.; Crump, T.; and Far, B. 2021. Reinforcement learning based recommender systems: A survey. *arXiv preprint arXiv:2101.06286*.
- Bagnell, J. A.; and Schneider, J. G. 2001. Autonomous helicopter control using reinforcement learning policy search methods. In *ICRA*. IEEE.
- Barto, A.; Mirolli, M.; and Baldassarre, G. 2013. Novelty or surprise? *Frontiers in psychology*, 4: 907.
- Bellemare, M. G.; Ostrovski, G.; Guez, A.; Thomas, P.; and Munos, R. 2016a. Increasing the action gap: New operators for reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30.
- Bellemare, M. G.; Srinivasan, S.; Ostrovski, G.; Schaul, T.; Saxton, D.; and Munos, R. 2016b. Unifying Count-Based Exploration and Intrinsic Motivation. *Conference on Neural Information Processing Systems (NIPS)*.
- Benbrahim, H.; and Franklin, J. A. 1997. Biped dynamic walking using reinforcement learning. *Robotics and Autonomous Systems*, 22(3-4): 283–302.
- Brafman, R. I.; and Tennenholtz, M. 2002. R-max-a general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research*, 3(Oct): 213–231.
- Buckman, J.; Gelada, C.; and Bellemare, M. G. 2021. The importance of pessimism in fixed-dataset policy optimization. *International Conference on Learning Representations (ICLR)*.
- Burda, Y.; Edwards, H.; Pathak, D.; Storkey, A.; Darrell, T.; and Efros, A. A. 2018. Large-scale study of curiosity-driven learning. *arXiv preprint arXiv:1808.04355*.
- Burda, Y.; Edwards, H.; Storkey, A.; and Klimov, O. 2019. Exploration by random network distillation. *International Conference on Learning Representations*.
- Dadashi, R.; Rezaeifar, S.; Vieillard, N.; Hussenot, L.; Pietquin, O.; and Geist, M. 2021. Offline Reinforcement Learning with Pseudometric Learning. *Self-Supervision for Reinforcement Learning Workshop - ICLR 2021*.
- Endo, G.; Morimoto, J.; Matsubara, T.; Nakanishi, J.; and Cheng, G. 2008. Learning CPG-based biped locomotion with a policy gradient method: Application to a humanoid robot. *The International Journal of Robotics Research*, 27(2): 213–228.
- Ernst, D.; Geurts, P.; and Wehenkel, L. 2005. Tree-based batch mode reinforcement learning. *Journal of Machine Learning Research*, 6: 503–556.
- Fu, J.; Kumar, A.; Nachum, O.; Tucker, G.; and Levine, S. 2020. D4rl: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*.
- Fujimoto, S.; Hoof, H.; and Meger, D. 2018. Addressing function approximation error in actor-critic methods. In *International Conference on Machine Learning*, 1587–1596. PMLR.
- Fujimoto, S.; Meger, D.; and Precup, D. 2019. Off-policy deep reinforcement learning without exploration. In *International Conference on Machine Learning*.
- Geist, M.; Scherrer, B.; and Pietquin, O. 2019. A theory of regularized markov decision processes. In *International Conference on Machine Learning*, 2160–2169. PMLR.
- Ghasemipour, S. K. S.; Schuurmans, D.; and Gu, S. S. 2020. Emaq: Expected-max q-learning operator for simple yet effective offline and online rl. *arXiv preprint arXiv:2007.11091*.
- Ghavamzadeh, M.; Petrik, M.; and Chow, Y. 2016. Safe policy improvement by minimizing robust baseline regret. *Advances in Neural Information Processing Systems*.
- Grigorescu, S.; Trasnea, B.; Cocias, T.; and Macesanu, G. 2020. A survey of deep learning techniques for autonomous driving. *Journal of Field Robotics*, 37(3): 362–386.
- Gulcehre, C.; Colmenarejo, S. G.; Wang, Z.; Sygnowski, J.; Paine, T.; Zolna, K.; Chen, Y.; Hoffman, M.; Pascanu, R.; and de Freitas, N. 2021. Regularized Behavior Value Estimation. *arXiv preprint arXiv:2103.09575*.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. In *International Conference on Machine Learning (ICML)*.
- Henderson, P.; Islam, R.; Bachman, P.; Pineau, J.; Precup, D.; and Meger, D. 2018. Deep Reinforcement Learning that Matters. *AAAI Conference on Artificial Intelligence*.
- Jin, Y.; Yang, Z.; and Wang, Z. 2021. Is Pessimism Provably Efficient for Offline RL? *arXiv:2012.15085*.
- Kingma, D. P.; and Welling, M. 2014. Auto-Encoding Variational Bayes. In *Proceedings of the International Conference on Learning Representations*.
- Kober, J.; Bagnell, J. A.; and Peters, J. 2013. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11): 1238–1274.
- Kumar, A.; Fu, J.; Tucker, G.; and Levine, S. 2019. Stabilizing off-policy q-learning via bootstrapping error reduction. *Neural Information Processing Systems (NeurIPS)*.
- Kumar, A.; Zhou, A.; Tucker, G.; and Levine, S. 2020. Conservative q-learning for offline reinforcement learning. *Neural Information Processing Systems (NeurIPS)*.
- Lagoudakis, M. G.; and Parr, R. 2003. Least-squares policy iteration. *The Journal of Machine Learning Research*.
- Laroche, R.; Trichelair, P.; and Des Combes, R. T. 2019. Safe policy improvement with baseline bootstrapping. In *International Conference on Machine Learning*, 3652–3661. PMLR.
- Levine, S.; Kumar, A.; Tucker, G.; and Fu, J. 2020. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2016. Continuous

- control with deep reinforcement learning. In *International Conference on Learning Representations (ICLR)*.
- Mandel, T.; Liu, Y.-E.; Levine, S.; Brunskill, E.; and Popovic, Z. 2014. Offline policy evaluation across representations with applications to educational games. In *AAMAS*, 1077–1084.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; and Riedmiller, M. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.
- Murphy, S. A.; van der Laan, M. J.; Robins, J. M.; and Group, C. P. P. R. 2001. Marginal mean models for dynamic regimes. *Journal of the American Statistical Association*, 96(456): 1410–1423.
- Nair, A.; Dalal, M.; Gupta, A.; and Levine, S. 2020. Accelerating online reinforcement learning with offline datasets. *arXiv preprint arXiv:2006.09359*.
- Ostrovski, G.; Bellemare, M. G.; Oord, A.; and Munos, R. 2017. Count-based exploration with neural density models. In *International conference on machine learning*, 2721–2730. PMLR.
- Pathak, D.; Agrawal, P.; Efros, A. A.; and Darrell, T. 2017. Curiosity-driven exploration by self-supervised prediction. In *International Conference on Machine Learning*, 2778–2787. PMLR.
- Peng, X. B.; Kumar, A.; Zhang, G.; and Levine, S. 2019. Advantage-weighted regression: Simple and scalable off-policy reinforcement learning. *arXiv preprint arXiv:1910.00177*.
- Petrik, M.; and Subramanian, D. 2014. RAAM: The benefits of robustness in approximating aggregated MDPs in reinforcement learning. *Advances in Neural Information Processing Systems*, 27: 1979–1987.
- Pidhorskyi, S.; Almoheisen, R.; Adjeroh, D. A.; and Doretto, G. 2018. Generative probabilistic novelty detection with adversarial autoencoders. *Advances in neural information processing systems*.
- Puterman, M. L. 2014. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.
- Rajeswaran, A.; Kumar, V.; Gupta, A.; Vezzani, G.; Schulman, J.; Todorov, E.; and Levine, S. 2017. Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. *Proceedings of Robotics: Science and Systems (RSS)*.
- Rashidinejad, P.; Zhu, B.; Ma, C.; Jiao, J.; and Russell, S. 2021. Bridging offline reinforcement learning and imitation learning: A tale of pessimism. *arXiv preprint arXiv:2103.12021*.
- Rezende, D. J.; Mohamed, S.; and Wierstra, D. 2014. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32, 1278–1286.
- Rojanavas, P.; Srinil, P.; and Pinngern, O. 2005. New recommendation system using reinforcement learning. *Special Issue of the Intl. J. Computer, the Internet and Management*, 13(SP 3).
- Sallab, A. E.; Abdou, M.; Perot, E.; and Yogamani, S. 2017. Deep reinforcement learning framework for autonomous driving. *Electronic Imaging*, 2017(19): 70–76.
- Schmidhuber, J. 1991. A Possibility for Implementing Curiosity and Boredom in Model-Building Neural Controllers. In *Proceedings of the First International Conference on Simulation of Adaptive Behavior on From Animals to Animals*, 222–227. Cambridge, MA, USA: MIT Press. ISBN 0262631385.
- Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; et al. 2018. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419): 1140–1144.
- Sohn, K.; Lee, H.; and Yan, X. 2015. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems*, 28: 3483–3491.
- Tang, H.; Houthoofd, R.; Foote, D.; Stooke, A.; Chen, X.; Duan, Y.; Schulman, J.; De Turck, F.; and Abbeel, P. 2017. # exploration: A study of count-based exploration for deep reinforcement learning. In *31st Conference on Neural Information Processing Systems (NIPS)*, volume 30, 1–18.
- Thrun, S. 1995. An approach to learning mobile robot navigation. *Robotics and Autonomous systems*, 15(4): 301–319.
- Vieillard, N.; Kozuno, T.; Scherrer, B.; Pietquin, O.; Munos, R.; and Geist, M. 2020. Leverage the average: an analysis of KL regularization in reinforcement learning. In *NeurIPS-34th Conference on Neural Information Processing Systems*.
- Vieillard, N.; Pietquin, O.; and Geist, M. 2020. Munchausen Reinforcement Learning. *Advances in Neural Information Processing Systems*, 33.
- Wang, R.; Ciliberto, C.; Amadori, P.; and Demiris, Y. 2019. Random Expert Distillation: Imitation Learning via Expert Policy Support Estimation. *arXiv:1905.06750*.
- Wang, Z.; Novikov, A.; Żoła, K.; Springenberg, J. T.; Reed, S.; Shahriari, B.; Siegel, N.; Merel, J.; Gulcehre, C.; Heess, N.; et al. 2020. Critic regularized regression. *Neural Information Processing Systems (NeurIPS)*.
- Wu, Y.; Tucker, G.; and Nachum, O. 2019. Behavior regularized offline reinforcement learning. *arXiv preprint arXiv:1911.11361*.
- Xu, D.; Ricci, E.; Yan, Y.; Song, J.; and Sebe, N. 2015. Learning deep representations of appearance and motion for anomalous event detection. *Computer Vision and Image Understanding*.
- Yu, T.; Kumar, A.; Rafailov, R.; Rajeswaran, A.; Levine, S.; and Finn, C. 2021. Combo: Conservative offline model-based policy optimization. *arXiv preprint arXiv:2102.08363*.
- Yu, T.; Thomas, G.; Yu, L.; Ermon, S.; Zou, J.; Levine, S.; Finn, C.; and Ma, T. 2020. Mopo: Model-based offline policy optimization. *Neural Information Processing Systems (NeurIPS)*.
- Zhou, C.; and Paffenroth, R. C. 2017. Anomaly detection with robust deep autoencoders. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, 665–674.