

Offline Reinforcement Learning with Mixture of Deterministic Policies

Anonymous authors

Paper under double-blind review

Abstract

Offline reinforcement learning (RL) has recently attracted considerable attention as an approach for utilizing past experiences to learn a policy. Recent studies have reported the challenges of offline RL, such as estimating the values of actions that are outside the data distribution. To mitigate offline RL issues, we propose an algorithm that leverages a mixture of deterministic policies. When the data distribution is multimodal, fitting a policy modeled with a unimodal distribution, such as Gaussian distribution, may lead to interpolation between separate modes, thereby resulting in the value estimation of actions that are outside the data distribution. In our framework, the state-action space is divided by learning discrete latent variables, and the sub-policies corresponding to each region are trained. The proposed algorithm was derived by considering the variational lower bound of the offline RL objective function. We show empirically that the use of the proposed mixture policy can reduce the accumulation of the critic loss in offline RL, which was reported in previous studies. Experimental results also indicate that using a mixture of deterministic policies in offline RL improves the performance with the D4RL benchmarking datasets.

1 Introduction

Reinforcement learning (RL) (Sutton & Barto, 2018) has achieved remarkable success in various applications. Many of its successes have been achieved in online learning settings where the RL agent interacts with the environment during the learning process. However, such interactions are often time-consuming and computationally expensive. The aim of reducing the number of interactions in RL has spurred active interest in offline RL (Levine et al., 2020), also known as batch RL (Lange et al., 2012). In offline RL, the goal is to learn the optimal policy from a prepared dataset generated through arbitrary and unknown processes. Prior work on offline RL has focused on how to avoid estimating the Q-values of actions that are outside the data distribution (Fujimoto et al., 2019; Fujimoto & Gu, 2021). In this study, we propose addressing it from the perspective of the policy structure. Our hypothesis is that the evaluation of the out-of-distribution actions can be avoided by dividing the state-action space, which can be achieved by learning discrete latent variables of the state-action space. When the data distribution is multimodal, as shown in Figure 1(a), fitting a policy modeled with a unimodal distribution, such as Gaussian distribution, may lead to interpolation between separate modes, thereby resulting in the value estimation of actions that are outside the data distribution (Figure 1(b)). To avoid this, we employ a mixture of deterministic policies (Figure 1(c)). We divide the state-action space and learn the sub-policies for each region. Ideally, this approach can help avoid interpolating separate modes of the data distribution.

The main contributions of this study are as follows: 1) it provides a practical algorithm for training a mixture of deterministic policies in offline RL and 2) investigates the effect of policy structure in offline RL. Although it is expected that a mixture of deterministic policies has advantages over a monolithic policy, it is not trivial to train a mixture of deterministic policies. We derived the proposed algorithm by considering the variational lower bound of the offline RL objective function. We refer to the proposed algorithm as deterministic mixture policy optimization (DMPO). Additionally, we proposed a regularization technique for a mixture policy based on mutual information. We empirically demonstrate that the proposed regularization technique improves the performance of the proposed algorithm. A previous study (Brandfonbrener et al.,

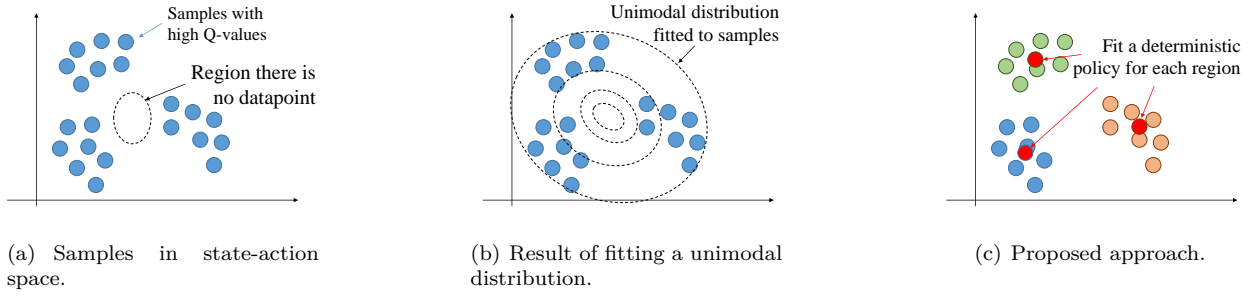


Figure 1: Illustration of the proposed approach. (a) In offline RL, the distribution of samples is often multimodal; (b) Fitting a unimodal distribution to such samples can lead to generating the action out of the data distribution; (c) In the proposed approach, first the latent discrete variable of the state-action space is learned, and then a deterministic policy is learned for each region.

2021) reported the accumulation of critic loss values during the training phase, which was attributed to generating out-of-distribution actions. In our experiments, we investigated the effect of the policy structures in offline RL through comparison with methods that use a monolithic deterministic policy, Gaussian policy, and Gaussian mixture policy. We empirically show that the use of a mixture of deterministic policies can reduce the accumulation of the approximation error in offline RL. Although a mixture of Gaussian policies has been used in the online RL literature, we show that the use of a Gaussian mixture policy does not significantly improve the performance of an offline RL algorithm. Through experiments with benchmark tasks in D4RL (Fu et al., 2020), we demonstrate that the proposed algorithms are competitive with prevalent offline RL methods.

2 Related Work

Recent studies have shown that regularization is a crucial component of offline RL (Fujimoto et al., 2019; Kumar et al., 2020; Levine et al., 2020; Kostrikov et al., 2021). For example, Kostrikov et al. proposed a regularization based on Fischer divergence Kostrikov et al. (2021), and Fujimoto et al. showed that simply adding a behavior cloning term to the objective function in TD3 can achieve state-of-the-art performance on D4RL benchmark tasks Fujimoto & Gu (2021). Other studies have investigated the structure of the critic, proposing the use of an ensemble of critics (An et al., 2021) or offering a one-step offline RL approach (Brandfonbrener et al., 2021; Goo & Niekum, 2021). Previous studies (Fujimoto et al., 2019; Fujimoto & Gu, 2021) have indicated that the source of the value approximation error is “extrapolation error” that occurs when the value of state-action pairs that are not contained in a given dataset is estimated. We hypothesize that such an “extrapolation error” can be mitigated by dividing the state-action space, which can be potentially achieved by learning discrete latent variables. We investigate the effect of incorporating policy structure as an inductive bias in offline RL, which has not been thoroughly investigated.

Learning the discrete latent variable in the context of RL is closely related to a mixture policy, where a policy is represented as a combination of a finite number of sub-policies. In a mixture policy, one of the sub-policies is activated for a given state, and the module that determines which sub-policy is to be used is often called the gating policy (Daniel et al., 2016). Because of its two-layered structure, a mixture policy is also called a *hierarchical policy* (Daniel et al., 2016). Although we did not consider temporal abstraction in this study, we note that a well-known hierarchical RL framework with temporal abstraction is the option critic (Bacon et al., 2017). Because we consider policies without temporal abstraction, we use the term “mixture policy,” following the terminology in Wulfmeier et al. (2021). Previous studies have demonstrated the advantages of mixture policies for online RL (Osa et al., 2019; Zhang & Whiteson, 2019; Wulfmeier et al., 2020; 2021; Akrouer et al., 2021; Sasaki & Matsubara, 2021). In these existing methods, sub-policies are often trained to cover separate modes of the Q-function, which is similar to our approach. Although existing methods have leveraged the latent variable in offline RL (Zhou et al., 2020; Chen et al., 2021b; 2022), the latent variable is continuous in these methods. For example, Chen et al. recently proposed an algorithm called latent-variable advantage-weighted policy optimization (LAPO), which leverages continuous

latent space for policy learning Chen et al. (2022). LAPO incorporates an importance weight based on the advantage function and learns the continuous latent variable. Although LAPO can achieve state-of-the-art performance on well-known benchmark tasks, we empirically show in this study that LAPO suffers from a surge of the critic loss during training.

3 Problem Formulation

Reinforcement Learning Consider a reinforcement learning problem under a Markov decision process (MDP) defined by a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, r, \gamma, d)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, $\mathcal{P}(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t)$ is the transition probability density, $r(\mathbf{s}, \mathbf{a})$ is the reward function, γ is the discount factor, and $d(\mathbf{s}_0)$ is the probability density of the initial state. A policy $\pi(\mathbf{a}|\mathbf{s}) : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$ is defined as the conditional probability density over the actions given the states. The goal of RL is to identify a policy that maximizes the expected return $\mathbb{E}[R_0|\pi]$, where the return is the sum of the discounted rewards over time given by $R_t = \sum_{k=t}^T \gamma^{k-t} r(\mathbf{s}_k, \mathbf{a}_k)$. The Q-function, $Q^\pi(\mathbf{s}, \mathbf{a})$, is defined as the expected return when starting from state \mathbf{s} and taking action \mathbf{a} , then following policy π under a given MDP (Sutton & Barto, 2018).

In offline RL, it is assumed that the learning agent is provided with a fixed dataset, $\mathcal{D} = \{(\mathbf{s}_i, \mathbf{a}_i, r_i)\}_{i=1}^N$, comprising states, actions, and rewards collected by an unknown behavior policy. The goal of offline RL is to obtain a policy that maximizes the expected return using \mathcal{D} without online interactions with the environment during the learning process.

Objective function We formulate the offline RL problem as follows: given dataset $\mathcal{D} = \{(\mathbf{s}_i, \mathbf{a}_i, r_i)\}_{i=1}^N$ obtained through the interactions between behavior policy $\beta(\mathbf{a}|\mathbf{s})$ and the environment, our goal is to obtain policy π that maximizes the expected return. In the process of training a policy in offline RL, the expected return is evaluated with respect to the states stored in the given dataset. Thus, the objective function is given by:

$$J(\pi) = \mathbb{E}_{\mathbf{s} \sim \mathcal{D}, \mathbf{a} \sim \pi} [f^\pi(\mathbf{s}, \mathbf{a})], \quad (1)$$

where f^π is a function that quantifies the performance of policy π . There are several choices for f^π as indicated in Schulman et al. (2016). TD3 employed the action-value function, $f^\pi(\mathbf{s}, \mathbf{a}) = Q^\pi(\mathbf{s}, \mathbf{a})$, and A2C employed the advantage-function $f^\pi(\mathbf{s}, \mathbf{a}) = A^\pi(\mathbf{s}, \mathbf{a})$ (Mnih et al., 2016). Other previous studies employed shaping with an exponential function, such as $f^\pi(\mathbf{s}, \mathbf{a}) = \exp(Q^\pi(\mathbf{s}, \mathbf{a}))$ (Peters & Schaal, 2007) or $f^\pi(\mathbf{s}, \mathbf{a}) = \exp(A^\pi(\mathbf{s}, \mathbf{a}))$ (Neumann & Peters, 2008; Wang et al., 2018). Without a loss of generality, we assume that the objective function is given by equation 1. We derive the proposed algorithm by considering the lower bound of the objective function of offline RL in equation 1.

Mixture policy In this study, we consider a mixture of policies given by

$$\pi(\mathbf{a}|\mathbf{s}) = \sum_{\mathbf{z} \in \mathcal{Z}} \pi_{\text{gate}}(\mathbf{z}|\mathbf{s}) \pi_{\text{sub}}(\mathbf{a}|\mathbf{s}, \mathbf{z}), \quad (2)$$

where \mathbf{z} is a discrete latent variable, $\pi_{\text{gate}}(\mathbf{z}|\mathbf{s})$ is the gating policy that determines the value of the latent variable, and $\pi_{\text{sub}}(\mathbf{a}|\mathbf{s}, \mathbf{z})$ is the sub-policy that determines the action for a given \mathbf{s} and \mathbf{z} . We assume that a sub-policy $\pi_{\text{sub}}(\mathbf{a}|\mathbf{s}, \mathbf{z})$ is deterministic; the sub-policy determines the action for a given \mathbf{s} and \mathbf{z} in a deterministic manner as $\mathbf{a} = \boldsymbol{\mu}_\theta(\mathbf{s}, \mathbf{z})$, where $\boldsymbol{\mu}_\theta(\mathbf{s}, \mathbf{z})$ is parameterized by vector $\boldsymbol{\theta}$. Additionally, we assume that the gating policy $\pi_{\text{gate}}(\mathbf{z}|\mathbf{s})$ determines the latent variable as:

$$\mathbf{z} = \arg \max_{\mathbf{z}'} Q_w(\mathbf{s}, \boldsymbol{\mu}_\theta(\mathbf{s}, \mathbf{z}')), \quad (3)$$

where $Q_w(\mathbf{s}, \mathbf{a})$ is the estimated Q-function parameterized by vector \mathbf{w} . This gating policy is applicable to objective functions such as $f^\pi(\mathbf{s}, \mathbf{a}) = \exp(Q^\pi(\mathbf{s}, \mathbf{a}))$, $f^\pi(\mathbf{s}, \mathbf{a}) = A^\pi(\mathbf{s}, \mathbf{a})$, and $f^\pi(\mathbf{s}, \mathbf{a}) = \exp(A^\pi(\mathbf{s}, \mathbf{a}))$. Please refer to A for details.

4 Training a mixture of deterministic policies by maximizing the variational lower bound

We consider a training procedure based on policy iteration (Sutton & Barto, 2018), in which the critic and policy are iteratively improved. In this section, we describe the policy update procedure of the proposed method.

4.1 Variational lower bound for offline RL

To derive the update rule for policy parameter θ , we first consider the lower bound of objective function $\log J(\pi)$ in equation 1. We assume that $f^\pi(\mathbf{s}, \mathbf{a})$ in equation 1 is approximated with $\hat{f}_w^\pi(\mathbf{s}, \mathbf{a})$, which is parameterized with a vector w . In a manner similar to Dayan & Hinton (1997); Kober & Peters (2011), when $\hat{f}_w^\pi(\mathbf{s}, \mathbf{a}) > 0$ for any \mathbf{s} and \mathbf{a} , we can determine the lower bound of $\log J(\pi)$ using Jensen’s inequality as follows:

$$\log J(\pi) \approx \log \int d^\beta(\mathbf{s}) \pi_\theta(\mathbf{a}|\mathbf{s}) \hat{f}_w^\pi(\mathbf{s}, \mathbf{a}) d\mathbf{s} d\mathbf{a} \quad (4)$$

$$= \log \int d^\beta(\mathbf{s}) \beta(\mathbf{a}|\mathbf{s}) \frac{\pi_\theta(\mathbf{a}|\mathbf{s})}{\beta(\mathbf{a}|\mathbf{s})} \hat{f}_w^\pi(\mathbf{s}, \mathbf{a}) d\mathbf{s} d\mathbf{a} \quad (5)$$

$$\geq \int d^\beta(\mathbf{s}) \beta(\mathbf{a}|\mathbf{s}) \log \frac{\pi_\theta(\mathbf{a}|\mathbf{s})}{\beta(\mathbf{a}|\mathbf{s})} \hat{f}_w^\pi(\mathbf{s}, \mathbf{a}) d\mathbf{s} d\mathbf{a} \quad (6)$$

$$= \mathbb{E}_{(\mathbf{s}, \mathbf{a}) \sim \mathcal{D}} \left[\log \pi_\theta(\mathbf{a}|\mathbf{s}) \hat{f}_w^\pi(\mathbf{s}, \mathbf{a}) \right] - \mathbb{E}_{(\mathbf{s}, \mathbf{a}) \sim \mathcal{D}} \left[\log \beta(\mathbf{a}|\mathbf{s}) \hat{f}_w^\pi(\mathbf{s}, \mathbf{a}) \right], \quad (7)$$

where $\beta(\mathbf{a}|\mathbf{s})$ is the behavior policy used for collecting the given dataset, and $d^\beta(\mathbf{s})$ is the stationary distribution over the state induced by executing behavior policy $\beta(\mathbf{a}|\mathbf{s})$. The second term in equation 7 is independent of policy parameter θ . Thus, we can maximize the lower bound of $J(\pi)$ by maximizing $\sum_{i=1}^N \log \pi_\theta(\mathbf{a}_i|\mathbf{s}_i) \hat{f}_w^\pi(\mathbf{s}_i, \mathbf{a}_i)$. When we employ $f^\pi(\mathbf{s}, \mathbf{a}) = \exp(A^\pi(\mathbf{s}, \mathbf{a}))$, and the policy is Gaussian, the resulting algorithm is equivalent to AWAC (Nair et al., 2020). To employ a mixture policy with a discrete latent variable, we further analyze the objective function in equation 7. As in Kingma & Welling (2014); Sohn et al. (2015), we obtain a variant of the variational lower bound of the conditional log-likelihood:

$$\begin{aligned} \log \pi_\theta(\mathbf{a}_i|\mathbf{s}_i) &\geq -D_{\text{KL}}(q_\phi(\mathbf{z}|\mathbf{s}_i, \mathbf{a}_i) || p(\mathbf{z}|\mathbf{s}_i)) + \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z}|\mathbf{s}_i, \mathbf{a}_i)} [\log \pi_\theta(\mathbf{a}_i|\mathbf{s}_i, \mathbf{z})] \\ &= \ell_{\text{cvae}}(\mathbf{s}_i, \mathbf{a}_i; \theta, \phi), \end{aligned} \quad (8)$$

where $q_\phi(\mathbf{z}|\mathbf{s}, \mathbf{a})$ is the approximate posterior distribution parameterized with vector ϕ , and $p(\mathbf{z}|\mathbf{s})$ is the true posterior distribution. The derivation of equation 8 is provided in B. Although it is often assumed in prior studies (Fujimoto et al., 2019) that \mathbf{z} is statistically independent of \mathbf{s} , that is, $p(\mathbf{z}|\mathbf{s}) = p(\mathbf{z})$, in our framework, $p(\mathbf{z}|\mathbf{s})$ should represent the behavior of the gating policy, $\pi_\theta(\mathbf{z}|\mathbf{s})$. Recognizing the challenge of exactly representing the gating policy $\pi_\theta(\mathbf{z}|\mathbf{s})$ in equation 3, we approximate it using the Softmax distribution given by

$$p(\mathbf{z}|\mathbf{s}) = \frac{\exp(Q_w(\mathbf{s}, \boldsymbol{\mu}_\theta(\mathbf{s}, \mathbf{z})))}{\sum_{\mathbf{z} \in \mathcal{Z}} \exp(Q_w(\mathbf{s}, \boldsymbol{\mu}_\theta(\mathbf{s}, \mathbf{z})))}. \quad (9)$$

Since we employ double-clipped Q-learning as in Fujimoto et al. (2018), we compute

$$Q_w(\mathbf{s}, \boldsymbol{\mu}_\theta(\mathbf{s}, \mathbf{z})) = \min_{j=1,2} Q_{w_j}(\mathbf{s}, \boldsymbol{\mu}_\theta(\mathbf{s}, \mathbf{z})) \quad (10)$$

in our implementation. The second term in equation 8 is approximated as the mean squared error, similar to that in the standard implementation of VAE. Based on equation 7 and equation 8, we obtain the objective function for training the mixture policy as follows:

$$\mathcal{L}_{\text{ML}}(\theta, \phi) = \sum_{i=1}^N f^\pi(\mathbf{s}_i, \mathbf{a}_i) \ell_{\text{cvae}}(\mathbf{s}_i, \mathbf{a}_i; \theta, \phi). \quad (11)$$

This objective can be regarded as the weighted maximum likelihood (Kober & Peters, 2011) of a mixture policy. Our objective function can be viewed as reconstructing the state-action pairs with adaptive weights, similar to that in Peters & Schaal (2007); Nair et al. (2020). Therefore, the policy samples actions within the support and does not evaluate out-of-distribution actions. The primary difference between the proposed and existing methods (Peters & Schaal, 2007; Nair et al., 2020) is that: the use of a mixture of policies conditioned on discrete latent variables in our approach can be regarded as dividing the state-action space. For example, in AWAC (Nair et al., 2020), a unimodal policy was used to reconstruct all of the “good” actions in the given dataset. However, in the context of offline RL, the given dataset may contain samples collected by diverse behaviors and enforcing the policy to cover all modes in the dataset can degrade the resulting performance. In our approach, policy $\pi_{\theta}(\mathbf{a}|\mathbf{s}, \mathbf{z})$ is encouraged to mimic the state-action pairs that are assigned to the same values of \mathbf{z} without mimicking the actions that are assigned to different values of \mathbf{z} .

4.2 Mutual-information-based regularization

To improve the performance of the mixture of deterministic policies, we propose a regularization technique for a mixture policy based on the mutual information (MI) between \mathbf{z} and the state action pair (\mathbf{s}, \mathbf{a}) , which we denote by $I(\mathbf{z}; \mathbf{s}, \mathbf{a})$. As shown in Barber & Agakov (2003), the variational lower bound of $I(\mathbf{z}; \mathbf{s}, \mathbf{a})$ is given as follows:

$$\begin{aligned} I(\mathbf{s}, \mathbf{a}; \mathbf{z}) &= H(\mathbf{z}) - H(\mathbf{z}|\mathbf{s}, \mathbf{a}) = \mathbb{E}_{(\mathbf{s}, \mathbf{a}, \mathbf{z}) \sim p_{\pi}} [\log p(\mathbf{z}|\mathbf{s}, \mathbf{a})] + H(\mathbf{z}) \\ &= \mathbb{E}_{(\mathbf{s}, \mathbf{a}) \sim \beta(\mathbf{s}, \mathbf{a})} [D_{\text{KL}}(p(\mathbf{z}|\mathbf{s}, \mathbf{a}) || g_{\psi}(\mathbf{z}|\mathbf{s}, \mathbf{a}))] + \mathbb{E}_{(\mathbf{s}, \mathbf{a}, \mathbf{z}) \sim p} [\log g_{\psi}(\mathbf{z}|\mathbf{s}, \mathbf{a})] + H(\mathbf{z}) \\ &\geq \mathbb{E}_{(\mathbf{s}, \mathbf{a}, \mathbf{z}) \sim p} [\log g_{\psi}(\mathbf{z}|\mathbf{s}, \mathbf{a})] + H(\mathbf{z}), \end{aligned} \quad (12)$$

where $g_{\psi}(\mathbf{z}|\mathbf{s}, \mathbf{a})$ is an auxiliary distribution to approximate the posterior distribution $p(\mathbf{z}|\mathbf{s}, \mathbf{a})$.

Thus, the final objective function is as follows:

$$\mathcal{L}(\theta, \phi, \psi) = \mathcal{L}_{\text{ML}}(\theta, \phi) + \lambda \sum_{i=1}^N \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})} \log g_{\psi}(\mathbf{z}|\mathbf{s}_i, \mu_{\theta}(\mathbf{s}_i, \mathbf{z})). \quad (13)$$

MI-based regularization using the second term in equation 13 encourages the diversity of the behaviors encoded in the sub-policy $\pi(\mathbf{a}|\mathbf{s}, \mathbf{z})$. In Section 7, we empirically show that this regularization improves the performance of the proposed method.

To implement MI-based regularization, we introduced a network to represent $g_{\psi}(\mathbf{z}|\mathbf{s}, \mathbf{a})$ in addition to a network that represents the posterior distribution $q_{\phi}(\mathbf{z}|\mathbf{s}, \mathbf{a})$. While maximizing the objective \mathcal{L}_{ML} in equation 11, both the actor $\mu_{\theta}(\mathbf{s}, \mathbf{z})$ and auxiliary distribution $g_{\psi}(\mathbf{z}|\mathbf{s}, \mathbf{a})$ are updated; however, the posterior distribution $q_{\phi}(\mathbf{z}|\mathbf{s}, \mathbf{a})$ is frozen. While maximizing $\sum_{i=1}^N \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})} \log g_{\psi}(\mathbf{z}|\mathbf{s}_i, \mu_{\theta}(\mathbf{s}_i, \mathbf{z}))$, both actor $\mu_{\theta}(\mathbf{s}, \mathbf{z})$ and posterior distribution $q_{\phi}(\mathbf{z}|\mathbf{s}, \mathbf{a})$ are updated, but the auxiliary distribution $g_{\psi}(\mathbf{z}|\mathbf{s}, \mathbf{a})$ is frozen. To maximize $\log g_{\psi}(\mathbf{z}|\mathbf{s}_i, \mu_{\theta}(\mathbf{s}_i, \mathbf{z}))$, the latent variable is sampled from the prior distribution, that is, the uniform distribution in this case, and the maximization of $\log g_{\psi}(\mathbf{z}|\mathbf{s}_i, \mu_{\theta}(\mathbf{s}_i, \mathbf{z}))$ is approximated by minimizing the squared difference, $\|\mathbf{z} - \hat{\mathbf{z}}\|^2$, where $\hat{\mathbf{z}}$ is the output of the network that represents $g_{\psi}(\mathbf{z}|\mathbf{s}, \mathbf{a})$.

5 Training the critic for a mixture of deterministic policies

To derive the objective function for training the critic for a mixture of deterministic policies using the gating policy in equation 3, we consider the following operator:

$$\mathcal{T}_{\mathbf{z}} Q_{\mathbf{z}} = r(\mathbf{s}, \mathbf{a}) + \gamma \mathbb{E}_{\mathbf{s}'} \left[\max_{\mathbf{z}'} Q_{\mathbf{z}}(\mathbf{s}', \mu(\mathbf{s}', \mathbf{z}')) \right]. \quad (14)$$

We refer to operator $\mathcal{T}_{\mathbf{z}}$ as the *latent-max-Q operator*. Following the method in Ghasemipour et al. (2021), we prove the following theorems.

Theorem 5.1. *In the tabular setting, $\mathcal{T}_{\mathbf{z}}$ is a contraction operator in the \mathcal{L}_{∞} norm. Hence, with repeated applications of $\mathcal{T}_{\mathbf{z}}$, any initial Q function converges to a unique fixed point.*

Algorithm 1 Deterministic mixture policy optimization (DMPO)

```

Initialize the actor  $\mu_\theta$ , critic  $Q_{w_j}$  for  $j = 1, 2$ , and the posterior  $q_\phi(\mathbf{z}|\mathbf{s}, \mathbf{a})$ 
for  $t = 1$  to  $T$  do
  Sample a minibatch  $\{(\mathbf{s}_i, \mathbf{a}_i, \mathbf{s}'_i, r_i)\}_{i=1}^M$  from  $\mathcal{D}$ 
  for each element  $(\mathbf{s}_i, \mathbf{a}_i, \mathbf{s}'_i, r_i)$  do
    Compute the target value as
     $y_i = r + \gamma \max_{\mathbf{z}' \in \mathcal{Z}} \min_{j=1,2} Q_{w'_j}(\mathbf{s}'_i, \mu_{\theta'}(\mathbf{s}'_i, \mathbf{z}'))$ 
  end for
  Update the critic by minimizing  $\sum_{i=1}^M \|y_i - Q_{w_j}(\mathbf{s}_i, \mathbf{a}_i)\|^2$  for  $j = 1, 2$ 
  if  $t \bmod d_{\text{interval}} = 0$  then
    Update the actor and the posterior by maximizing equation 11
    (optionally) Update the actor by maximizing  $\sum_{i=1}^M \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})} \log g_\psi(\mathbf{z}|\mathbf{s}_i, \mu_\theta(\mathbf{s}_i, \mathbf{z}))$ 
  end if
end for

```

The proof of Theorem 5.1 is provided in C.

Theorem 5.2. *Let $Q_{\mathbf{z}}$ denote the unique fixed point achieved in Theorem 5.1 and $\pi_{\mathbf{z}}$ denote the policy that chooses the latent variable as $\mathbf{z} = \arg \max_{\mathbf{z}'} Q(\mathbf{s}, \mu(\mathbf{s}, \mathbf{z}'))$ and outputs the action given by $\mu(\mathbf{s}, \mathbf{z})$ in a deterministic manner. Then $Q_{\mathbf{z}}$ is the Q-value function corresponding to $\pi_{\mathbf{z}}$.*

Proof. (Theorem 5.2) Rearranging equation 14 with $\mathbf{z}' = \arg \max Q_{\mathbf{z}}(\mathbf{s}', \mu(\mathbf{s}', \mathbf{z}'))$, we obtain

$$\mathcal{T}_{\mathbf{z}} Q_{\mathbf{z}} = r(\mathbf{s}, \mathbf{a}) + \gamma \mathbb{E}_{\mathbf{s}'} \mathbb{E}_{\mathbf{a}' \sim \pi_{\mathbf{z}}} [Q_{\mathbf{z}}(\mathbf{s}', \mathbf{a}')]. \quad (15)$$

Because $Q_{\mathbf{z}}$ is the unique fixed point of $\mathcal{T}_{\mathbf{z}}$, we have our result. \square

These theorems reveal that the latent-max-Q operator, $\mathcal{T}_{\mathbf{z}}$, retains the contraction and fixed-point existence properties. Based on these results, we estimate the Q-function by applying the latent-max-Q operator. In our implementation, we employed double-clipped Q-learning (Fujimoto et al., 2018). Thus, given dataset \mathcal{D} , the critic is trained by minimizing

$$\mathcal{L}(\mathbf{w}_j) = \sum_{(\mathbf{s}_i, \mathbf{a}_i, \mathbf{s}'_i, r_i) \in \mathcal{D}} \|Q_{w_j}(\mathbf{s}_i, \mathbf{a}_i) - y_i\|^2 \quad (16)$$

for $j = 1, 2$, where target value y_i is computed as

$$y_i = r_i + \gamma \max_{\mathbf{z}' \in \mathcal{Z}} \min_{j=1,2} Q_{w'_j}(\mathbf{s}'_i, \mu_{\theta'}(\mathbf{s}'_i, \mathbf{z}')). \quad (17)$$

6 Practical implementation

The proposed DMPO algorithm is summarized as Algorithm 1. Similar to that in TD3 (Fujimoto et al., 2018), the actor is updated once after d_{interval} updates of the critics. In our implementation, we set $d_{\text{interval}} = 2$. For modeling the discrete latent variable, we used the Gumbel-Softmax trick (Jang et al., 2017; Maddison et al., 2017). Additionally, we employed the state normalization method used in TD3+BC (Fujimoto & Gu, 2021). During preliminary experiments, we found that when $f^\pi(\mathbf{s}, \mathbf{a}) = \exp(bA^\pi(\mathbf{s}, \mathbf{a}))$ in equation 11, scaling factor b has non-trivial effects on performance and the best value of b differs for each task. To avoid changing the scaling parameter for each task, we used the normalization of the advantage function as

$$f^\pi(\mathbf{s}, \mathbf{a}) = \exp \left(\frac{\alpha (A^\pi(\mathbf{s}, \mathbf{a}) - \max_{(\tilde{\mathbf{s}}, \tilde{\mathbf{a}}) \in \mathcal{D}_{\text{batch}}} A^\pi(\tilde{\mathbf{s}}, \tilde{\mathbf{a}}))}{\max_{(\tilde{\mathbf{s}}, \tilde{\mathbf{a}}) \in \mathcal{D}_{\text{batch}}} A^\pi(\tilde{\mathbf{s}}, \tilde{\mathbf{a}}) - \min_{(\tilde{\mathbf{s}}, \tilde{\mathbf{a}}) \in \mathcal{D}_{\text{batch}}} A^\pi(\tilde{\mathbf{s}}, \tilde{\mathbf{a}})} \right), \quad (18)$$

where $\mathcal{D}_{\text{batch}}$ is a mini-batch sampled from the given dataset \mathcal{D} and α is a constant. We set $\alpha = 10$ for mujoco tasks in our experiments.




Figure 2: Performance on a simple task with multimodal data distribution.

Table 1: Algorithm setup in the experiment.

	TD3+BC	AWAC	DMPO
Critic training	double-clipped Q-learning	double-clipped Q-learning	double-clipped Q-learning
Policy type	monolithic & deterministic	monolithic & stochastic	mixture & deterministic
Regularization	BC term	none	none
State	normalized	normalized	normalized

7 Experiments

We investigated the effect of policy structure on the resulting performance and training errors of critics. In the first experiment, we performed a comparative assessment of TD3+BC (Fujimoto & Gu, 2021), AWAC (Nair et al., 2020) and DMPO on a toy problem where the distribution of samples in a given dataset is multimodal. Further, we conducted a quantitative comparison between the proposed and baseline methods with D4RL benchmark tasks (Fu et al., 2020). In the following section, we refer to the proposed method based on the objective in equation 11 as DMPO, and a variant of the proposed method with MI-based regularization in equation 13 as infoDMPO. In both, the toy problem and the D4RL tasks, we used the author-provided implementations of TD3+BC, and our implementations of DMPO and AWAC are based on the author-provided implementation of TD3+BC.

7.1 Multimodal data distribution on toy task

To show the effect of multimodal data distribution in a given dataset, we evaluated the performance of TD3+BC, AWAC, and DMPO on a toy task, as shown in Figure 2. The differences between the compared methods are summarized in Table 1. In our implementation of AWAC and DMPO, we used state normalization and double-clipped Q-learning as in TD3+BC and the normalization of the advantage function described in Section 6. The difference between AWAC and DMPO indicates the effect of the policy representation. In this toy task, the agent is represented as a point mass, the state is the position of the point mass in two-dimensional space, and the action is the small displacement of the point mass. There are two goals in this task, which are indicated by red circles in Figure 2. The blue circle denotes the starting position in Figure 2, and there are three obstacles, which are indicated by solid black circles. In this task, the reward is sparse: when the agent reaches one of the goals, it receives a reward of 1.0, and the episode ends. If the agent makes contact with one of the obstacles, the agent receives a reward of -1.0, and the episode ends. In the given dataset, trajectories for the two goals are provided, and there is no information on which goal the agent is heading to.

The scores are summarized in Table 2. Among the evaluated methods, only DMPO successfully solved this task. The policy obtained by TD3+BC did not reach its goal in a stable manner, as shown in Figure 2(b). Similarly, as shown in Figure 2(c), the policy learned by AWAC often slows down around point (0,0) and fails to reach the goal. This behavior implies that AWAC attempts to average over multiple modes of the distribution. In contrast, the policy learned by DMPO successfully reaches one of the goals. Because the main difference between AWAC and DMPO is the policy architecture, the result

Table 2: Performance on the toy task.

TD3+BC	AWAC	DMPO
-0.2 ± 0.4	0.33 ± 0.44	1.0 ± 0.0

shows that the unimodal policy distribution fails to deal with the multimodal data distribution, whereas the mixture policy employed in DMPO successfully dealt with it. The activation of the sub-policies is visualized in Figure 2(e). The color indicates the value of the discrete latent variable given by the gating policy, $\mathbf{z}^* = \arg \max_{\mathbf{z}} Q_w(\mathbf{s}, \boldsymbol{\mu}(\mathbf{s}, \mathbf{z}))$. Figure 2(d) shows that different sub-policies are activated for different regions, thereby indicating that DMPO appropriately divides the state-action space.

7.2 Effect of policy structure

We investigated the effect of policy structure by comparing the proposed method with existing methods that incorporate the importance weight based on the advantage function. We used AWAC and LAPO Chen et al. (2022) as baseline methods. To investigate the difference between the mixture of the stochastic policies and the mixture of the deterministic policies, we evaluated a variant of AWAC with Gaussian mixture policies, which we refer to as mixAWAC. All baseline methods used double-clipped Q-learning for the critic in this experiment.

We found that the authors’ implementation of LAPO¹ includes techniques to improve performance, such as action normalization and clipping of the target value for the state-value function. Although such techniques are compatible with DMPO and our baseline methods, they are not incorporated in our experiments for simplicity. We refer to LAPO without these techniques as LAPO-. LAPO- can be considered as a baseline method that incorporates a continuous latent variable in its policy structure.

The implementations of AWAC and DMPO were identical to those used in the previous experiment. In our evaluation, $|Z| = 8$ was used for DMPO and infoDMPO. D presents the effect of the dimensionality of the discrete latent variables. In this study, we evaluated the baseline methods with mujoco-v2 and antmaze-v0 tasks.

7.2.1 Performance score on D4RL

A comparison between AWAC, mixAWAC, LAPO-, and DMPO is presented in Table 3. These methods incorporate importance weights based on the advantage function with different policy structures. Therefore, the differences between these methods indicate the effect of policy structure. In our experiments, we did not observe significant differences in the performance of AWAC and mixAWAC. This result indicates that the use of Gaussian mixture policies does not lead to performance improvement. However, the performance of DMPO matched or exceeded that of AWAC, except for the Hopper-expert and Hopper-medium-expert tasks. This result also confirms that the use of a mixture of deterministic policies is beneficial for these tasks, although the benefits would be task-dependent.

The difference between mixAWAC and DMPO implies a difference between a Gaussian mixture policy and a mixture of deterministic policy. In a Gaussian mixture policy, there is a possibility that one of the Gaussian components covers a large action space and interpolates the separate modes of action. If this happens, out-of-distribution actions will be generated by the learned policy. However, in a mixture of the deterministic policy, there is no such possibility that one of the components covers a large action space.

In addition, DMPO outperformed LAPO-, which is the state-of-the-art method in offline RL. While LAPO- learns the continuous latent space based on advantage-weighting, DMPO learns the discrete latent variable using advantage-weighting. This result also indicates that the use of a discrete latent variable is beneficial for offline RL tasks. A comparison with additional baseline methods is provided in E.

7.2.2 Critic loss function

To investigate the effect of the policy structure on the critic loss function, we compared the value of the critic loss function among AWAC, mixAWAC, LAPO-, and DMPO. The normalized scores and value of the critic loss function during training are depicted in Figure 3. The value of the critic loss given by equation 16 is plotted for every 5,000 updates. Previous studies have indicated that the critic loss value can accumulate over iterations (Brandfonbrener et al., 2021). Figure 3 shows the accumulation of the critic loss in AWAC

¹<https://github.com/pcchenxi/LAPO-offlienRL>

Table 3: Comparison with methods incorporating advantage-weighting using D4RL-v2 datasets. Average normalized scores over the last 10 test episodes and five seeds are shown. The boldface text indicates the best performance. HC = HalfCheetah, HP = Hopper, WK = Walker2d.

		AWAC	mixAWAC	LAPO-	DMPO
Expert	HC	94.8±0.2	94.0±0.5	94.9±0.2	97.0±1.0
	HP	109.8±2.9	111.8±0.8	111.4±0.6	93.6±15.1
	WK	111.0±0.2	110.5±0.3	111.1±0.2	111.4±0.3
Med.-E	HC	92.7±0.8	92.1±0.6	94.2±0.4	91.1±3.4
	HP	98.6±10.7	102.0±17.5	110.9±1.0	78.4±19.0
	WK	109.2±0.3	109.1±0.3	110.4±0.1	109.9±0.4
Med.-R	HC	40.9±0.6	41.5±0.4	40.7±1.9	45.2±0.8
	HP	38.2±9.4	41.2±4.7	52.6±11.3	89.2±8.1
	WK	65.0±15.7	67.7±8.8	60.1±11.7	82.1±3.8
Med.	HC	44.3±0.2	45.1±0.3	45.5±0.2	47.5±0.4
	HP	57.5±3.0	57.2±3.9	51.8±3.6	71.2±6.5
	WK	81.0±2.5	78.7 ± 4.8	77.7±3.5	79.4±4.7
Rand.	HC	3.2±1.3	2.2±0.0	2.3±0.3	15.8±1.6
	HP	7.3±0.9	8.2±0.2	21.8±11.3	12.0±10.0
	WK	3.1±1.0	4.9±1.1	5.6±3.1	2.5±2.6
Antmaze	umaze	49.8±6.2	57.4±6.2	76.4±3.3	83.6±4.5
	umaze-d.	53.8±13.0	46.8±6.9	56.6±8.6	43.2±7.8
	med.-p.	0.0±0.0	0.0±0.0	0.8±1.6	77.0±5.1
	med.-d.	0.0±0.0	0.0±0.0	0.0±0.0	56.8±27.2
	large-p.	0.0±0.0	0.0±0.0	0.0±0.0	1.0±1.3
	large-d.	0.0±0.0	0.0±0.0	0.0±0.0	4.8±9.6

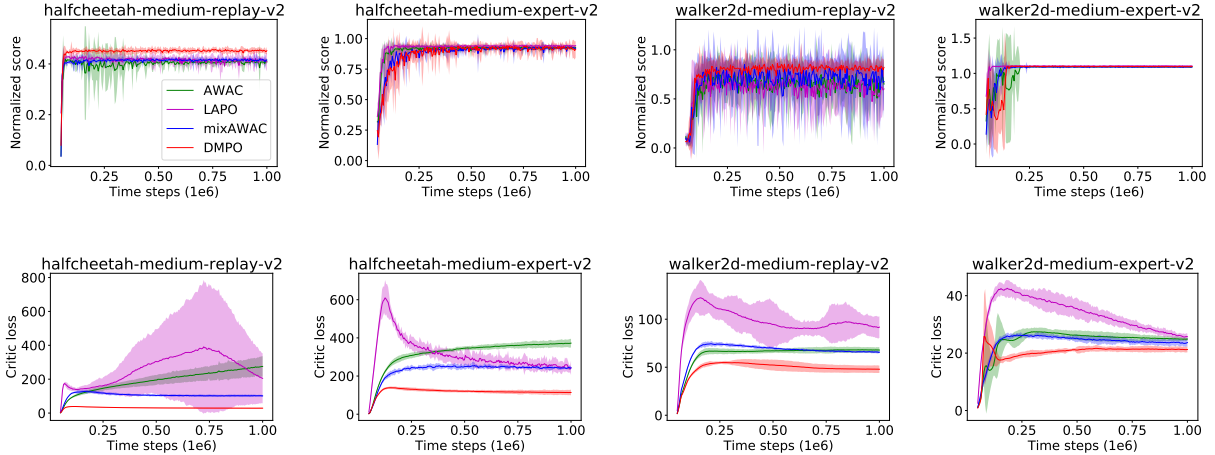


Figure 3: Critic loss and normalized score during the training.

on mujoco-v2 tasks. The difference between AWAC and mixAWAC indicates that using a Gaussian mixture policy often reduces the accumulation of the critic loss. The critic loss of mixAWAC is lower than that of AWAC in halfcheetah-medium-replay-v2 and halfcheetah-medium-expert-v2 tasks. This result shows that the use of a multimodal policy can reduce the accumulation of the critic loss in offline RL. In addition, the critic loss of DMPO is even lower than that of mixAWAC, and the result demonstrates that using a mixture of deterministic policies can further reduce the critic loss than using a Gaussian mixture policy. These results indicate that using a mixture of deterministic policies can reduce the generation of out-of-distribution actions, which is essential for offline RL.

Regarding LAPO-, the critic loss value increased rapidly at the beginning of the training. Although the critic loss value often decreases at the end of the LAPO- training, the critic loss value is still higher than that of DMPO. The surge in the critic loss value indicates the generation of out-of-distribution actions during training in LAPO-. Importantly, in DMPO, the value of the critic loss is clearly lower, and the performance of the policy is better than that of LAPO-. This result indicates that the use of a discrete latent variable can be more effective than using a continuous latent variable on these tasks. In Brandfonbrener et al. (2021), it was shown that the accumulation of critic loss values can be reduced by introducing regularization. Our




Figure 4: Histogram of the Bellman errors after 20k steps.

results indicate that the use of a mixture policy can also mitigate the accumulation of critic loss in offline RL, which suggests the importance of incorporating inductive bias in the policy structure. However, it is worth noting that the reduction in the critic loss given by equation 16 does not necessarily improve the policy performance. In halfcheetah-medium-expert-v2, although the critic loss was significantly lower in DMPO than in AWAC, there was no significant difference in performance between DMPO and AWAC. Recently, Fujimoto et al. indicated that a lower value of the critic loss given by equation 16 does not necessarily mean better performance, and the observation in Fujimoto et al. (2022) aligns with our experiments. The metric to measure the accuracy of the value estimation is still an open problem in RL.

As another qualitative result, Figure 4 shows the histograms of the Bellman error after training with 20 thousand steps. The Bellman error in mixAWAC is distributed more widely than that in AWAC, indicating that the use of the Gaussian mixture policy can increase the variance during the training of the critic. In contrast, the distribution of the Bellman error in DMPO is more narrow than those of AWAC, mixAWAC, and LAPO-, indicating that the use of the mixture of deterministic policies may lead to reducing the variance during the critic training. This variance reduction could also be considered a reason why DMPO outperformed baseline methods.

7.3 Comparison with prevalent baselines

We compared the performance of the proposed method with that of prevalent baselines. As baseline methods, we used TD3+BC, CQL (Kumar et al., 2020), IQL (Kostrikov et al., 2022), and LAPO. CQL incorporates conservative critic update and the entropy regularization. In the experiments reported in this section, we used the authors’ implementation of LAPO. IQL employs expectile regression for learning the critic to address the issue of generating out-of-distribution actions during training. Because the aim of our study is to investigate the policy structure, the approach of IQL, which address the critic learning, is orthogonal to ours. IQL is the state-of-the-art method for antmaze task on D4RL, which involves dealing with long horizons and requires “stitching” together sub-trajectories in a given dataset (Fu et al., 2020). In the implementation of IQL, several techniques, such as scheduling of the learning rate, were used to improve the performance. To compete with IQL on antmaze task, we also used techniques used in Chen et al. (2022). Therefore, the implementations of DMPO and infoDMPO for antmaze tasks are slightly different from those for other tasks. In our preliminary experiment, we evaluated IQL using the techniques proposed in Chen et al. (2022), and observed that the original implementation of IQL showed better performance. Therefore, we used with the original implementation of IQL for comparison. In this experiment, we used the mujoco-v2, antmaze-v0, and adroit tasks on D4RL.

A comparison of TD3+BC, CQL, and IQL is presented in Tables 4, 5, and 6. The boldface text indicates the best performance. In mujoco-v2 tasks, the performance of DMPO is comparable/superior to that of the state-of-the-art methods. In addition, infoDMPO, which employs MI-based regularization, outperformed DMPO on various tasks, and infoDMPO showed the best performance for 10 tasks among 15 mujoco-v2 tasks. This result shows that encouraging the diversity of sub-policies using the proposed MI-based regularization is effective for DMPO.

The advantages of DMPO and infoDMPO over TD3+BC and CQL are apparent for antmaze tasks. TD3+BC and CQL did not work satisfactorily on antmaze tasks, indicating that techniques used in these algorithms are not effective for such tasks. The performance of DMPO(ant ver.) and infoDMPO(ant ver.) on antmaze

Table 4: Results on mujoco tasks using D4RL-v2 datasets and AntMaze tasks. Average normalized scores over the last 10 test episodes and five seeds are shown. HC = HalfCheetah, HP = Hopper, WK = Walker2d.

		TD3+BC (re-run)	CQL (re-run)	IQL (re-run)	LAPO (re-run)	DMPO (ours)	infoDMPO (ours)
Expert	HC	96.3±0.9	22.0±9.6	96.1±1.5	95.4±0.3	97.0±1.0	95.6±2.0
	HP	109.9±2.5	105.8±3.8	98.4±13.1	110.9±2.3	93.6±15.1	107.5±2.9
	WK	110.2±0.4	108.9±0.4	112.6±0.3	111.5±0.2	111.4±0.3	112.1±0.4
Med.-E	HC	89.4±7.2	38.4±8.4	90.7±4.3	94.3±1.1	91.1±3.4	91.4±2.5
	HP	95.5±9.4	88.4±15.9	73.9±32.6	110.5±1.2	78.4±19.0	94.5±14.9
	WK	110.2±0.3	109.2±1.9	111.4±1.1	111.0±0.2	109.9±0.4	110.1±0.7
Med.-R	HC	44.7±0.4	46.9±0.3	43.6±1.4	41.9±1.0	45.2±0.8	46.7±0.6
	HP	73.8±18.9	95.5±1.7	90.6±14.3	59.7±14.2	89.2±8.1	98.5±2.0
	WK	64.5±17.0	77.5±3.1	82.2±3.6	50.3±18.6	82.1±3.8	86.7±3.2
Med.	HC	48.2±0.3	48.2±0.4	48.2±0.2	45.7±0.3	47.5±0.4	48.6±0.4
	HP	61.0±4.2	77.4±4.0	61.2±3.5	56.2±5.1	71.2± 6.5	86.4±7.6
	WK	84.7±1.3	81.5±2.5	82.9±6.0	80.5±1.8	79.4±4.7	85.0±0.8
Rand.	HC	11.5±0.6	24.1±1.5	12.6±4.6	27.1±1.0	15.8±1.6	16.3±1.2
	HP	8.7±0.3	2.2±1.9	7.4±0.3	15.2±8.6	12.0±10.0	20.4±9.8
	WK	1.4±1.9	4.3±7.9	5.5±1.6	2.2±1.5	2.5± 2.6	2.3±2.0
Total		1010.0	930.3	1017.3	1012.4	1026.5	1102.1

Table 5: Results on AntMaze tasks. Average normalized scores over the last 100 test episodes and five seeds are shown.

		TD3+BC (re-run)	CQL (re-run)	IQL (re-run)	LAPO (re-run)	DMPO (ant ver.) (ours)	infoDMPO (ant ver.) (ours)
Antmaze	umaze	92.8±2.7	73.0±4.9	87.4±4.5	97.2±2.7	92.8±2.1	89.4±5.1
	umaze-d.	45.0±22.2	43.8±4.4	64.6±5.6	57.4±11.7	32.6±25.6	34.8±18.0
	med.-p.	0.0±0.0	9.0±6.4	74.6±3.1	73.8±4.8	63.0±13.0	62.6±6.8
	med.-d.	0.0±0.0	3.8±4.2	73.8±7.1	81.0±3.6	75.0±8.5	82.8±4.4
	large-p.	0.0±0.0	0.0±0.0	39.0±7.2	27.6±13.3	42.2±23.0	47.4±14.5
	large-d.	0.0±0.0	0.0±0.4	48.0±9.0	26.2±17.5	56.6±4.5	38.0±4.8

tasks is comparable to that of IQL, which is the state-of-the-art method for these tasks. We observed similar results for adroit tasks. DMPO and infoDMPO clearly outperformed TD3+BC and CQL, and the performance of DMPO and infoDMPO was comparable to IQL. Considering that infoDMPO outperformed IQL on mujoco-v2 task, overall performance of infoDMPO is better than that of IQL. This result reveals that the use of a mixture of deterministic policies can result in a significant performance improvement in offline RL.

As a qualitative evaluation, we investigated the activation of sub-policies in DMPO. Activation of sub-policies in DMPO on the pen-human-v0 task is depicted in Figure 5. In Figure 5(a), the top row depicts the state at the 20th, 40th, 60th, and 80th time steps, and graphs in the middle row of the figure show the action-values of each sub-policy at each state, $Q_w(s, \mu(s, z))$. Figures 5 (a) and (b) show the results for different episodes. A previous study on the option framework (Smith et al., 2018) reported that in the existing method only a few options are activated and that the remaining options do not learn meaningful behaviors. In contrast, the results in Figure 5 show that the value of each of the sub-policies $Q_w(s, \mu(s, z))$ changes over time, and various sub-policies are activated during execution. This result implies the following: meaningful sub-policies are learned in DMPO and different behaviors are adaptively used to perform complicated manipulation task.

8 Conclusion

We presented DMPO, an algorithm for training a mixture of deterministic policies in offline RL. This algorithm can be interpreted as an approach that divides the state-action space by learning the discrete latent variable and the corresponding sub-policies in each region. In this study, we empirically investigated the effect of policy structure in offline RL. The experimental results reveal that the use of a mixture of

Table 6: Results on adroit tasks using the average normalized scores over the last 10 test episodes and five seeds.

		TD3+BC (re-run)	CQL(ρ) (re-run)	IQL (re-run)	LAPO (re-run)	DMPO (ours)	infoDMPO (ours)
Human	pen	0.8 \pm 8.0	98.3\pm81.8	88.8 \pm 21.2	78.9 \pm 14.1	86.1 \pm 8.8	94.8\pm16.5
	hammer	0.9 \pm 0.8	-7.1 \pm 0.1	1.0 \pm 0.2	1.1 \pm 0.4	1.2 \pm 0.2	2.4\pm0.9
	door	-0.3 \pm 0.0	-3.3 \pm 7.8	2.4 \pm 2.1	3.2 \pm 1.6	1.3 \pm 1.5	4.2\pm3.1
	relocate	-0.3 \pm 0.0	0.3\pm2.4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.1	0.1\pm0.0
Cloned	pen	0.5 \pm 7.0	-1.7 \pm 1.5	39.2 \pm 15.4	25.6 \pm 12.2	36.0 \pm 17.7	46.4\pm16.7
	hammer	0.2 \pm 0.0	-7.0 \pm 0.1	0.9 \pm 0.4	0.7 \pm 0.4	0.8 \pm 0.6	1.2\pm0.3
	door	-0.3 \pm 0.0	-9.4 \pm 0.0	0.6 \pm 1.2	0.7 \pm 0.9	0.0 \pm 0.0	0.8\pm0.8
	relocate	-0.3 \pm 0.0	-2.1 \pm 0.0	-0.2\pm0.0	-0.2 \pm 0.0	-0.2\pm0.0	-0.2\pm0.0

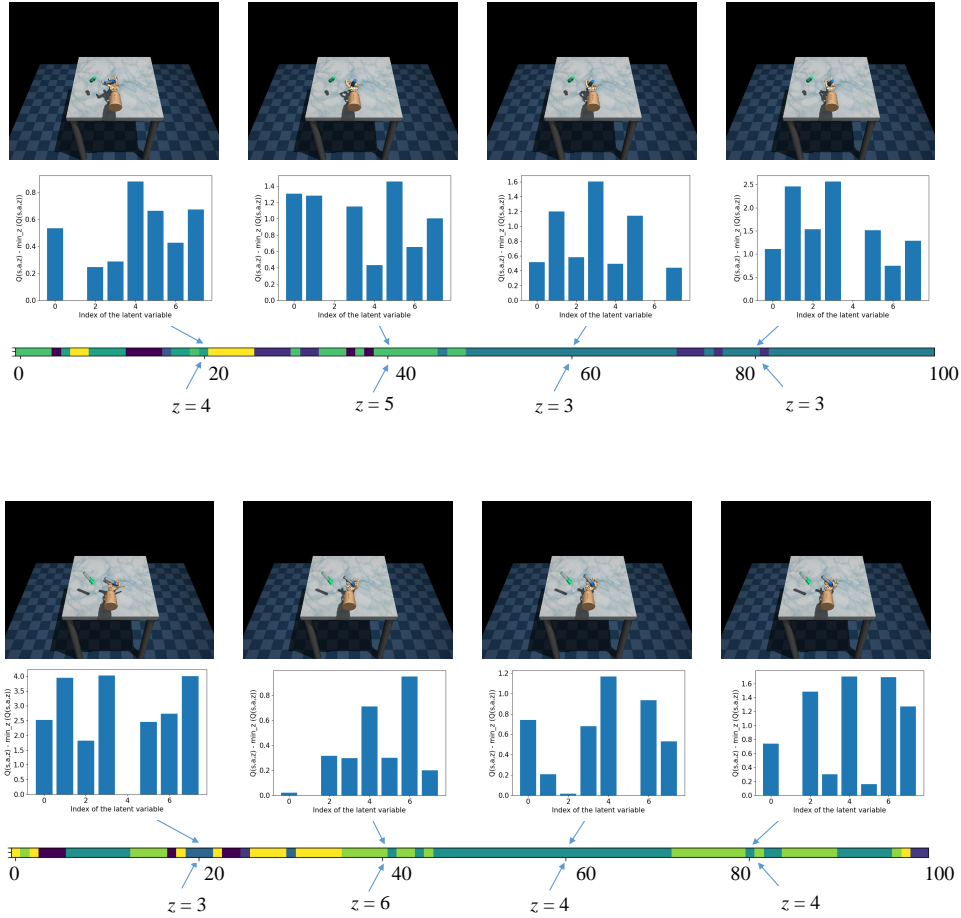


Figure 5: Visualization of sub-policy activation in the pen-human-v0 task. The top row depicts the state at the 20th, 40th, 60th, and 80th time steps; the graphs in the middle row depicts the action-values of each sub-policy at each state.

deterministic policies can mitigate the issue of critic error accumulation in offline RL. In addition, the results indicate that the use of a mixture of deterministic policies significantly improves the performance of an offline RL algorithm. We believe that our study contributes to advancing techniques to leverage policy structure in offline RL.

References

- Riad Akrou, Davide Tateo, and Jan Peters. Continuous action reinforcement learning from a mixture of interpretable experts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. doi: 10.1109/TPAMI.2021.3103132.
- Gaon An, Seungyong Moon, Jang-Hyun Kim, and Hyun Oh Song. Uncertainty-based offline reinforcement learning with diversified q-ensemble. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Pierre-Luc Bacon, Jean Harb, and Doina Precup. Option-critic architecture. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2017.
- David Barber and Felix Agakov. The im algorithm : A variational approach to information maximization. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2003.
- David Brandfonbrener, William F. Whitney, Rajesh Ranganath, and Joan Bruna. Offline rl without off-policy evaluation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Michael Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: Reinforcement learning via sequence modeling. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021a.
- Xi Chen, Ali Ghadirzadeh, Tianhe Yu, Jianhao Wang, Yuan Gao, Wenzhe Li, Bin Liang, Chelsea Finn, and Chongjie Zhang. Lapo: Latent-variable advantage-weighted policy optimization for offline reinforcement learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- Xiong-Hui Chen, Yang Yu, Qingyang Li, Fan-Ming Luo, Zhiwei Qin, Wenjie Shang, and Jieping Ye. Offline model-base adaptable policy learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021b.
- Christian Daniel, Gerhard Neumann, Oliver Kroemer, and Jan Peters. Hierarchical relative entropy policy search. *Journal of Machine Learning Research*, 17(93):1–50, 2016.
- P. Dayan and G. Hinton. Using expectation-maximization for reinforcement learning. *Neural Computation*, 9:271–278, 1997.
- Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. D4rl: Datasets for deep data-driven reinforcement learning. *arXiv*, 2020.
- Scott Fujimoto and Shixiang Shane Gu. A minimalist approach to offline reinforcement learning. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Scott Fujimoto, Herke van Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 1587–1596, 2018.
- Scott Fujimoto, David Meger, and Doina Precup. Off-policy deep reinforcement learning without exploration. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 2052–2062, 2019.
- Scott Fujimoto, David Meger, Doina Precup, Ofir Nachum, and Shixiang Shane Gu. Why should i trust you, bellman? the bellman error is a poor replacement for value error. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2022.
- Seyed Kamyar Seyed Ghasemipour, Dale Schuurmans, and Shixiang Shane Gu. Emaq: Expected-max q-learning operator for simple yet effective offline and online rl. In *Proceedings of the International Conference on Machine Learning (ICML)*, volume 139, pp. 3682–3691. PMLR, 18–24 Jul 2021.
- Wonjoon Goo and Scott Niekum. You only evaluate once: a simple baseline algorithm for offline rl. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2021.

- Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2014.
- J. Kober and J. Peters. Policy search for motor primitives in robotics. *Machine Learning*, 84:171–203, 2011.
- Ilya Kostrikov, Rob Fergus, Jonathan Tompson, and Ofir Nachum. Offline reinforcement learning with fisher divergence critic regularization. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2021.
- Ilya Kostrikov, Ashvin Nair, and Sergey Levine. Offline reinforcement learning with implicit q-learning. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2022.
- Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. Conservative q-learning for offline reinforcement learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Sascha Lange, Thomas Gabel, and Martin Riedmiller. Batch reinforcement learning. In *Reinforcement Learning*, pp. 45–73, 2012.
- Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv*, 2020.
- Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
- Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2016.
- Ashvin Nair, Abhishek Gupta, Murtaza Dalal, and Sergey Levine. AWAC: Accelerating online reinforcement learning with offline datasets. *arXiv*, arXiv:2006.09359, 2020.
- Gerhard Neumann and Jan Peters. Fitted q-iteration by advantage weighted regression. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2008.
- Takayuki Osa, Voot Tangkaratt, and Masashi Sugiyama. Hierarchical reinforcement learning via advantage-weighted information maximization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.
- Jan Peters and Stefan Schaal. Reinforcement learning by reward-weighted regression for operational space control. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2007.
- Hikaru Sasaki and Takamitsu Matsubara. Variational policy search using sparse gaussian process priors for learning multimodal optimal actions. *Neural Networks*, 143:291–302, 2021.
- John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2016.
- Matthew Smith, Herke Hoof, and Joelle Pineau. An inference-based policy gradient method for learning options. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2018.
- Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2015.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, Second edition, 2018.

E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033, 2012.

Qing Wang, Jiechao Xiong, Lei Han, peng sun, Han Liu, and Tong Zhang. Exponentially weighted imitation learning for batched historical data. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

Markus Wulfmeier, Abbas Abdolmaleki, Roland Hafner, Jost Tobias Springenberg, Michael Neunert, Tim Hertweck, Thomas Lampe, Noah Siegel, Nicolas Heess, and Martin Riedmiller. Compositional transfer in hierarchical reinforcement learning. In *Proceedings of Robotics: Science and Systems (R:SS)*, 2020.

Markus Wulfmeier, Dushyant Rao, Roland Hafner, Thomas Lampe, Abbas Abdolmaleki, Tim Hertweck, Michael Neunert, Dhruva Tirumala, Noah Siegel, Nicolas Heess, and Martin Riedmiller. Data-efficient hindsight off-policy option learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2021.

Shangdong Zhang and Shimon Whiteson. Dac: The double actor-critic architecture for learning options. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

Wenxuan Zhou, Sujay Bajracharya, and David Held. Plas: Latent action space for offline reinforcement learning. In *Proceedings of the Conference on Robot Learning (CoRL)*, volume 155, pp. 1719–1735, 2020.

A Applicability of the gating policy

In the proposed algorithm, we employ a gating policy that determines the value of the latent variable as follows:

$$z = \arg \max_{z'} Q_w(s, \mu_\theta(s, z')), \quad (19)$$

where $\mu_\theta(s, z')$ represents the deterministic sub-policy, and $Q_w(s, a)$ is the approximated Q-function. While this gating policy appears specific to the case where $Q^\pi(s, a)$ is maximized, it is applicable to other objective function such as $A^\pi(s, a)$, $\exp(Q^\pi(s, a))$, and $\exp(A^\pi(s, a))$. The advantage function is defined as $A^\pi(s, a) = Q^\pi(s, a) - V^\pi(s)$. Because the state value function $V^\pi(s)$ is independent of the action, we can obtain the following equation:

$$\arg \max_a Q^\pi(s, a) = \arg \max_a (Q^\pi(s, a) - V^\pi(s)) \quad (20)$$

$$= \arg \max_a A^\pi(s, a). \quad (21)$$

Thus, we can rewrite the gating policy as

$$z = \arg \max_{z'} Q_w(s, \mu_\theta(s, z')) \quad (22)$$

$$= \arg \max_{z'} A_w(s, \mu_\theta(s, z')). \quad (23)$$

Similarly, exponential function $\exp(\cdot)$ is a monotonically increasing function. Thus, the extrema of $Q^\pi(s, a)$ is also the extrema of $\exp(Q^\pi(s, a))$. Consequently, we can also rewrite the gating policy as

$$z = \arg \max_{z'} Q_w(s, \mu_\theta(s, z')) \quad (24)$$

$$= \arg \max_{z'} \exp(Q_w(s, \mu_\theta(s, z'))) \quad (25)$$

$$= \arg \max_{z'} A_w(s, \mu_\theta(s, z')) \quad (26)$$

$$= \arg \max_{z'} \exp(A_w(s, \mu_\theta(s, z'))). \quad (27)$$

Because we used this gating policy, it is deterministic in our implementation.

B Derivation of the variational lower bound

We employed the variational lower bound in equation 8 to derive the objective function for the proposed method. Here, we provide a detailed derivation, which was omitted in the main manuscript. We denote the true distribution induced by the policy $\pi_{\theta}(\mathbf{a}|\mathbf{s})$ as $p(\cdot)$, and the distribution that approximates the true distribution is denoted as $q(\cdot)$. The KL divergence between $q(\mathbf{x})$ and $p(\mathbf{x})$ is defined as

$$D_{\text{KL}}(q(\mathbf{x})||p(\mathbf{x})) = \int q(\mathbf{x}) \log \frac{q(\mathbf{x})}{p(\mathbf{x})} d\mathbf{z}. \quad (28)$$

Based on the above notation, the log-likelihood $\log \pi_{\theta}(\mathbf{a}_i|\mathbf{s}_i)$ can be written as follows:

$$\log \pi_{\theta}(\mathbf{a}_i|\mathbf{s}_i) = \int q_{\phi}(\mathbf{z}|\mathbf{s}_i, \mathbf{a}_i) \log \pi_{\theta}(\mathbf{a}_i|\mathbf{s}_i) d\mathbf{z} \quad (29)$$

$$= \int q_{\phi}(\mathbf{z}|\mathbf{s}_i, \mathbf{a}_i) (\log \pi(\mathbf{a}_i|\mathbf{s}_i, \mathbf{z}) + \log p(\mathbf{z}|\mathbf{s}_i) - \log p(\mathbf{z}|\mathbf{s}_i, \mathbf{a}_i)) d\mathbf{z} \quad (30)$$

$$\begin{aligned} &= \int q_{\phi}(\mathbf{z}|\mathbf{s}_i, \mathbf{a}_i) \log \frac{q(\mathbf{z}|\mathbf{s}_i, \mathbf{a}_i)}{p(\mathbf{z}|\mathbf{s}_i, \mathbf{a}_i)} d\mathbf{z} \\ &\quad - \int q(\mathbf{z}|\mathbf{s}_i, \mathbf{a}_i) \log \frac{q(\mathbf{z}|\mathbf{s}_i, \mathbf{a}_i)}{p(\mathbf{z}|\mathbf{s}_i)} d\mathbf{z} \\ &\quad + \int q(\mathbf{z}|\mathbf{s}_i, \mathbf{a}_i) \log \pi_{\theta}(\mathbf{a}_i|\mathbf{s}_i, \mathbf{z}) d\mathbf{z} \end{aligned} \quad (31)$$

$$\begin{aligned} &= D_{\text{KL}}(q(\mathbf{z}|\mathbf{s}_i, \mathbf{a}_i)||p(\mathbf{z}|\mathbf{s}_i, \mathbf{a}_i)) - D_{\text{KL}}(q(\mathbf{z}|\mathbf{s}_i, \mathbf{a}_i)||p(\mathbf{z}|\mathbf{s}_i)) \\ &\quad + \mathbb{E}_{\mathbf{z} \sim q(\mathbf{z}|\mathbf{s}_i, \mathbf{a}_i)} [\log \pi_{\theta}(\mathbf{a}_i|\mathbf{s}_i, \mathbf{z})]. \end{aligned} \quad (32)$$

In the first line, we consider marginalization over \mathbf{z} . As $\log \pi(\mathbf{a}|\mathbf{s})$ is independent of the latent variable \mathbf{z} , the equality in the first line holds. Because $D_{\text{KL}}(q(\mathbf{z}|\mathbf{s}, \mathbf{a})||p(\mathbf{z}|\mathbf{s}, \mathbf{a})) > 0$, we can obtain a variant of the variational lower bound of the conditional log-likelihood:

$$\log \pi_{\theta}(\mathbf{a}_i|\mathbf{s}_i) \geq -D_{\text{KL}}(q_{\phi}(\mathbf{z}|\mathbf{s}_i, \mathbf{a}_i)||p(\mathbf{z}|\mathbf{s}_i)) + \mathbb{E}_{\mathbf{z} \sim q(\mathbf{z}|\mathbf{s}_i, \mathbf{a}_i)} [\log \pi_{\theta}(\mathbf{a}_i|\mathbf{s}_i, \mathbf{z})]. \quad (33)$$

C Proof of contraction of the latent-max-Q operator

We consider operator $\mathcal{T}_{\mathbf{z}}$, which is given by

$$\mathcal{T}_{\mathbf{z}}Q(\mathbf{s}, \mathbf{a}) = \mathbb{E}_{\mathbf{s}'} \left[r(\mathbf{s}, \mathbf{a}) + \gamma \max_{\mathbf{z}} Q(\mathbf{s}', \boldsymbol{\mu}(\mathbf{s}', \mathbf{z}')) \right]. \quad (34)$$

To prove the contraction of $\mathcal{T}_{\mathbf{z}}$, we use the infinity norm given by

$$\|Q_1 - Q_2\|_{\infty} = \max_{\mathbf{s} \in \mathcal{S}, \mathbf{a} \in \mathcal{A}} |Q_1(\mathbf{s}, \mathbf{a}) - Q_2(\mathbf{s}, \mathbf{a})|, \quad (35)$$

where Q_1 and Q_2 are different estimates of the Q-function. We consider the infinity norm of the difference between the two estimates, Q_1 and Q_2 , after applying operator $\mathcal{T}_{\mathbf{z}}$:

$$\|\mathcal{T}_{\mathbf{z}}Q_1 - \mathcal{T}_{\mathbf{z}}Q_2\|_{\infty} \quad (36)$$

$$= \left| \mathbb{E}_{\mathbf{s}'} \left[r(\mathbf{s}, \mathbf{a}) + \gamma \max_{\mathbf{z}'} Q_1(\mathbf{s}', \boldsymbol{\mu}(\mathbf{s}', \mathbf{z}')) \right] - \mathbb{E}_{\mathbf{s}'} \left[r(\mathbf{s}, \mathbf{a}) + \gamma \max_{\mathbf{z}'} Q_2(\mathbf{s}', \boldsymbol{\mu}(\mathbf{s}', \mathbf{z}')) \right] \right| \quad (37)$$

$$= \left| \gamma \mathbb{E}_{\mathbf{s}'} \left[\max_{\mathbf{z}'} Q_1(\mathbf{s}', \boldsymbol{\mu}(\mathbf{s}', \mathbf{z}')) \right] - \gamma \mathbb{E}_{\mathbf{s}'} \left[\max_{\mathbf{z}'} Q_2(\mathbf{s}', \boldsymbol{\mu}(\mathbf{s}', \mathbf{z}')) \right] \right| \quad (38)$$

$$= \gamma \left| \mathbb{E}_{\mathbf{s}'} \left[\max_{\mathbf{z}'} Q_1(\mathbf{s}', \boldsymbol{\mu}(\mathbf{s}', \mathbf{z}')) \right] - \mathbb{E}_{\mathbf{s}'} \left[\max_{\mathbf{z}'} Q_2(\mathbf{s}', \boldsymbol{\mu}(\mathbf{s}', \mathbf{z}')) \right] \right| \quad (39)$$

$$= \gamma \left| \mathbb{E}_{\mathbf{s}'} \left[\max_{\mathbf{z}'} Q_1(\mathbf{s}', \boldsymbol{\mu}(\mathbf{s}', \mathbf{z}')) - \max_{\mathbf{z}'} Q_2(\mathbf{s}', \boldsymbol{\mu}(\mathbf{s}', \mathbf{z}')) \right] \right| \quad (40)$$

$$\leq \gamma \left| \mathbb{E}_{\mathbf{s}'} \|Q_1 - Q_2\|_{\infty} \right| \quad (41)$$

$$\leq \gamma \|Q_1 - Q_2\|_{\infty}. \quad (42)$$

The above relationship shows the contraction of operator $\mathcal{T}_{\mathbf{z}}$.

Table 7: Effect of the dimensionality of the discrete latent variable. WK=walker2d.

	infoDMPO $ Z = 4$	infoDMPO $ Z = 8$	infoDMPO $ Z = 16$	infoDMPO $ Z = 32$
pen-human-v0	75.7 \pm 18.9	94.8 \pm 16.5	75.0 \pm 17.5	86.7 \pm 12.4
WK-expert-v2	99.7 \pm 17.9	112.1\pm0.4	108.8 \pm 6.8	106.4\pm10.2
WK-med.-expert-v2	89.1 \pm 25.7	110.1\pm0.7	96.0 \pm 17.0	109.9\pm0.6
WK-med.-replay-v2	81.6 \pm 4.5	86.7\pm3.2	85.4 \pm 3.7	86.3\pm3.1
WK-med.-v2	81.8 \pm 2.5	85.0\pm0.8	69.9 \pm 28.3	84.3\pm1.0

Table 8: Results on mujoco tasks using D4RL-v2 datasets. Average normalized scores over the last 10 test episodes and five seeds are shown. HC = HalfCheetah, HP = Hopper, WK = Walker2d. The gray text indicates the performance lower than that of DMPO/infoDMPO. The bold text indicates the best performance.

		MAPLE (paper)	Decision Transformer (paper)	DMPO (ours)	infoDMPO (ours)
Med.- Expert	HC	63.5 \pm 6.5	86.8 \pm 1.3	91.1\pm 3.4	91.4\pm 2.5
	HP	42.5 \pm 4.1	107.6 \pm 1.8	78.4 \pm 19.0	94.5 \pm 14.9
	WK	73.8 \pm 8.0	108.1 \pm 0.2	109.9\pm 0.4	110.1\pm 0.7
Med.- Replay	HC	59.0 \pm 0.6	36.6 \pm 0.8	45.2 \pm 0.8	46.7 \pm 0.6
	HP	87.5 \pm 10.8	82.7 \pm 7.0	89.2 \pm 8.1	98.5\pm 2.0
	WK	76.7 \pm 3.8	66.6 \pm 3.0	82.1 \pm 3.8	86.7\pm 3.2
Med.	HC	50.4 \pm 1.9	42.6 \pm 0.1	47.5 \pm 0.4	48.6 \pm 0.4
	HP	21.1 \pm 1.2	67.6 \pm 1.0	71.2 \pm 6.5	86.4\pm 7.6
	WK	56.3 \pm 10.6	74.0 \pm 1.4	79.4 \pm 4.7	85.0\pm 0.8

D Effect of dimensionality of the discrete latent variable

In our evaluation, we first examined the effect of the dimensionality of the discrete latent variable. The results are presented in Table 7. As shown, infoDMPO with $|Z| = 8$ demonstrated the best performance, while the performance with $|Z| = 16$ and $|Z| = 32$ is comparable. These results show that the policy performance is not very sensitive to the dimensionality of the latent variable. However, the performance with $|Z| = 4$ is relatively weak, thereby indicating that the policy may not be sufficiently expressive when the dimensionality of the latent variable is significantly small. Because $|Z| = 8$ consistently provided satisfactory performance, $|Z| = 8$ was used in the subsequent evaluations.

E Comparison with additional baselines

We provide a comparison with additional baselines for the mujoco-v2 tasks in D4RL in Table 8. We present the results of MAPLE, which is a recent model-based offline algorithm that uses latent representations (Chen et al., 2021b). In addition, we provide the results of decision transformer (Chen et al., 2021a), as a representative transformer-based method. Although these methods are well-known and state-of-the-art, we focused on model-free and non-transformer-based methods in the main manuscript. For each baseline method, we adapted the results reported in the original paper. DMPO and infoDMPO provide a consistently better or comparable performance to these baseline methods, although our implementation of DMPO and infoDMPO does not employ techniques such as ensemble of critics. This result indicates a significant effect of the policy structure in offline RL.

F Hyperparameters and implementation details

Computational resource and license The experiments were run with Amazon Web Service and workstations with NVIDIA RTX 3090 GPUs and Intel Core i9-10980XE CPUs at 3.0GHz. We used the physics simulator, MuJoCo (Todorov et al., 2012) under an institutional license, and later we switched to the Apache license.

Software The software versions used in the experiments are listed below:

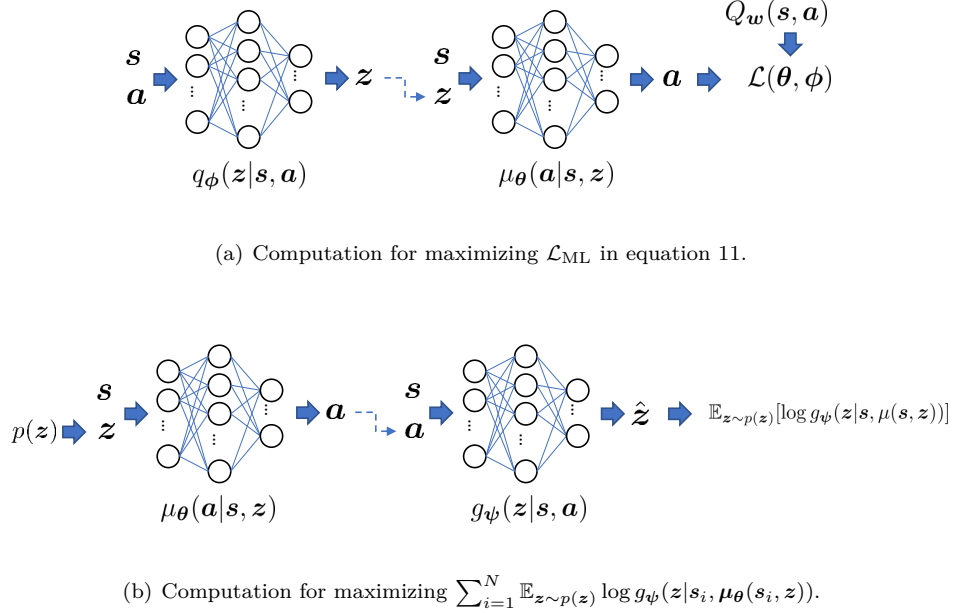


Figure 6: Connection between $q_\phi(z|s, a)$, $\mu_\theta(s, z)$, and $g_\psi(z|s, a)$ during training.

- Python 3.8
- Pytorch 1.10.0
- Gym 0.21.0
- MuJoCo 2.1.0
- mujoco-py 2.1.2.14

We used the author-provided implementations for TD3+BC² and CQL³. DMPO, AWAC, mixAWAC, and IQL were implemented based on the the author-provided implementation of TD3. For IQL, we used the hyperparameters provided in Kostrikov et al. (2022). To minimize the difference between DMPO and AWAC, we used a delayed update of the policy in both DMPO and AWAC. For simplicity, we did not use a regularization technique for the actor such as the dropout layer used in (Kostrikov et al., 2022), although the use of such techniques should further improve the performance. In our implementation of DMPO, the value of z is a part of the input to the actor network. Thus, different behaviors corresponding to different values of z are represented by the same actor network. The network architecture is illustrated in Figure 6.

Computation of the advantage function In DMPO, a policy is deterministic because both the gating policy $\pi(z|s)$ and sub-policy $\pi(a|s, z)$ are deterministic. Thus, the state-value function is given by

$$V^\pi(s) = \max_z Q^\pi(s, \mu(s, z)). \quad (43)$$

Therefore, the advantage function is given by

$$A^\pi(s, a) = Q^\pi(s, a) - V^\pi(s) = Q^\pi(s, a) - \max_z Q^\pi(s, \mu(s, z)). \quad (44)$$

²https://github.com/sfujim/TD3_BC

³<https://github.com/young-geng/CQL>

In the policy update, we use the target actor in the second term in equation 44. Thus, in our implementation, the advantage function is approximated as

$$A(\mathbf{s}, \mathbf{a}; \mathbf{w}, \boldsymbol{\theta}') = Q(\mathbf{s}, \mathbf{a}; \mathbf{w}) - \max_{\mathbf{z}} Q(\mathbf{s}, \boldsymbol{\mu}_{\boldsymbol{\theta}'}(\mathbf{s}, \mathbf{z}); \mathbf{w}). \quad (45)$$

Target smoothing in DMPO In DMPO, a policy is given by a mixture of deterministic sub-policies, where a sub-policy is selected in a deterministic manner, similar to that in equation 3. Thus, the mixture policy in this framework is deterministic. As reported in Fujimoto & Gu (2021), the use of a deterministic policy may lead to overfitting of the critic to narrow peaks. Because our policy is deterministic, we also employed a technique called target policy smoothing used in TD3. Thus, the target value in equation 17 is modified as follows:

$$y_i = r_i + \gamma \max_{\mathbf{z}' \in \mathcal{Z}} \min_{j=1,2} Q_{\mathbf{w}_j'}(\mathbf{s}', \boldsymbol{\mu}_{\boldsymbol{\theta}'}(\mathbf{s}', \mathbf{z}') + \epsilon_{\text{clip}}), \quad (46)$$

where ϵ_{clip} is given by

$$\epsilon_{\text{clip}} = \min(\max(\epsilon, -c), c) \quad \text{where} \quad \epsilon \sim \mathcal{N}(0, \sigma), \quad (47)$$

and constant c defines the range of the noise.

Techniques for Antmaze tasks In LAPO Chen et al. (2022), several techniques to stabilize the training of the value functions are used. Suppose the state-value function is approximated with $V_{\mathbf{w}_v}(\mathbf{s})$ parameterized with a vector \mathbf{w}_v , and the Q-function is approximated with two models, which are represented by $Q_{\mathbf{w}_j}(\mathbf{s}, \mathbf{a})$ for $j = 1, 2$. The state-value function is updated by minimizing

$$\mathcal{L}_v(\mathbf{w}_v) = \sum_{(\mathbf{s}_i, \mathbf{a}_i) \in \mathcal{D}} \|\tilde{y}_i - V_{\mathbf{w}_v}(\mathbf{s}_i)\|^2, \quad (48)$$

where the target value \tilde{y}_i is the clipped target value computed as

$$\tilde{y}_i = \max(\min(y_i, v_{\max}), v_{\min}) \quad (49)$$

and y_i is computed as

$$y_i = c \min_{j=1,2} Q_{\mathbf{w}_j}(\mathbf{s}_i, \mathbf{a}_i) + (1 - c) \max_{j=1,2} Q_{\mathbf{w}_j}(\mathbf{s}_i, \mathbf{a}_i) \quad (50)$$

and c is a constant, and we used $c = 0.7$ as in Chen et al. (2022). The minimum and maximum target value v_{\min} and v_{\max} are computed as

$$v_{\min} = \frac{1}{1 - \gamma} \min_{r_i \in \mathcal{D}} r_i \quad (51)$$

$$v_{\max} = \frac{1}{1 - \gamma} \max_{r_i \in \mathcal{D}} r_i. \quad (52)$$

The Q-function is updated by minimizing the following objective function:

$$\mathcal{L}_q(\mathbf{w}) = \sum_{(\mathbf{s}_i, \mathbf{a}_i, r_i, \mathbf{s}'_i) \in \mathcal{D}} \|r_i + \gamma V_{\mathbf{w}_v}(\mathbf{s}'_i) - Q_{\mathbf{w}}(\mathbf{s}_i, \mathbf{a}_i)\|^2. \quad (53)$$

For the antmaze tasks, we also used same techniques in DMPO.

Number of updates In the pen-human-v0, hammer-human-v0, door-human-v0, and relocate-human-v0 tasks, the number of samples contained in the dataset is significantly smaller than that for the other datasets. While the datasets for mujoco tasks contained approximately 1 million samples, the numbers of samples in the adroit-human tasks were as follows: pen-human-v0: 4,950 samples, hammer-human-v0: 11,264 samples, door-human-v0: 6,703 samples, and relocate-human-v0: 9,906 samples. Thus, in the pen-human-v0, hammer-human-v0, door-human-v0, and relocate-human-v0 tasks, we updated the policy 10,000 times, whereas for the other tasks, we updated the policy 1 million times. The aforementioned number of policy updates was applied to all methods.

Hyperparameters Tables 9–13 provide the hyperparameters used in the experiments.

Table 9: Hyperparameters of DMPO&infoDMPO.

	Hyperparameter	Value
Hyperparameters	Optimizer	Adam
	Critic learning rate	3e-4
	Actor learning rate	3e-4
	Posterior learning rate	3e-4
	Mini-batch size	256
	Discount factor	0.99
	Target update rate	5e-3
	Policy noise	0.2
	Policy noise clipping	(-0.5, 0.5)
	Policy update frequency	2
Architecture	Critic hidden dim	256
	Critic hidden layers	2
	Critic activation function	ReLU
	Actor hidden dim	256
	Actor hidden layers	2
	Actor activation function	ReLU
	Posterior hidden dim	256
	Posterior hidden layers	2
	Posterior activation function	ReLU
DMPO	Score scaling α	5.0 (human, Antmaze)
		10.0 (mujoco-v2)
infoDMPO	learning rate of the posterior for infomax	3e-6 (Adroit)
		5e-7 (mujoco-v2)
		5e-7 (Antmaze)
	Score scaling α	5.0 (Antmaze, HP-med.-expert)
		10.0 (others)

Table 10: Hyperparameters of AWAC.

	Hyperparameter	Value
Hyperparameters	Optimizer	Adam
	Critic learning rate	3e-4
	Actor learning rate	3e-4
	Mini-batch size	1024
	Discount factor	0.99
	Target update rate	5e-3
	Policy update frequency	2
	Score scaling α	10.0
Architecture	Critic hidden dim	256
	Critic hidden layers	2
	Critic activation function	ReLU
	Actor hidden dim	256
	Actor hidden layers	2
	Actor activation function	ReLU

Table 11: Hyperparameters of TD3+BC. The default hyperparameters in the TD3+BC GitHub are used.

	Hyperparameter	Value
Hyperparameters	Optimizer	Adam
	Critic learning rate	3e-4
	Actor learning rate	3e-4
	Mini-batch size	256
	Discount factor	0.99
	Target update rate	5e-3
	Policy noise	0.2
	Policy noise clipping	(-0.5, 0.5)
	Policy update frequency	2
	α	2.5
Architecture	Critic hidden dim	256
	Critic hidden layers	2
	Critic activation function	ReLU
	Actor hidden dim	256
	Actor hidden layers	2
	Actor activation function	ReLU

Table 12: Hyperparameters of CQL. The default hyperparameters in the CQL GitHub are used.

	Hyperparameter	Value
Hyperparameters	Optimizer	Adam
	Critic learning rate	3e-4
	Actor learning rate	3e-5
	Mini-batch size	256
	Discount factor	0.99
	Target update rate	5e-3
	Target entropy	-1·Action Dim
	Entropy in Q target	False
	Lagrange	False
	α	10
	Pre-training steps	40e3
	Num sampled actions (during eval)	10
	Num sampled actions (logsumexp)	10
Architecture	Critic hidden dim	256
	Critic hidden layers	3
	Critic activation function	ReLU
	Actor hidden dim	256
	Actor hidden layers	3
	Actor activation function	ReLU

Table 13: Hyperparameters of IQL. The default hyperparameters in the IQL in the paper Kostrikov et al. (2022) are used.

	Hyperparameter	Value
IQL hyperparameters	Optimizer	Adam
	Critic learning rate	3e-4
	Actor learning rate	3e-4
	Mini-batch size	256
	Discount factor	0.99
	Target update rate	5e-3
	Expectile	0.7 (mujoco-v2)
		0.7 (adroit)
		0.9 (antmaze)
	Advantage scale	3.0 (mujoco-v2)
Architecture		0.5 (adroit)
		10.0 (antmaze)
	Actor learning rate scheduling	cosine
	Critic hidden dim	256
	Critic hidden layers	2
	Critic activation function	ReLU
	Actor hidden dim	256
	Actor hidden layers	3
	Actor activation function	ReLU