EMAQ: EXPECTED-MAX Q-LEARNING OPERATOR FOR SIMPLE YET EFFECTIVE OFFLINE AND ONLINE RL

Anonymous authors
Paper under double-blind review

ABSTRACT

Off-policy reinforcement learning (RL) holds the promise of sample-efficient learning of decision-making policies by leveraging past experience. However, in the offline RL setting – where a fixed collection of interactions are provided and no further interactions are allowed – it has been shown that standard off-policy RL methods can significantly underperform. Recently proposed methods often aim to address this shortcoming by constraining learned policies to remain close to the given dataset of interactions. In this work, we closely investigate an important simplification of BCQ (Fujimoto et al., 2018a) – a prior approach for offline RL – which removes a heuristic design choice and naturally restrict extracted policies to remain *exactly* within the support of a given behavior policy. Importantly, in contrast to their original theoretical considerations, we derive this simplified algorithm through the introduction of a novel backup operator, Expected-Max Q-Learning (EMaQ), which is more closely related to the resulting practical algorithm. Specifically, in addition to the distribution support, EMaQ explicitly considers the number of samples and the proposal distribution, allowing us to derive new sub-optimality bounds which can serve as a novel measure of complexity for offline RL problems. In the offline RL setting – the main focus of this work – EMaQ matches and outperforms prior state-of-the-art in the D4RL benchmarks (Fu et al., 2020a). In the online RL setting, we demonstrate that EMaQ is competitive with Soft Actor Critic (SAC). The key contributions of our empirical findings are demonstrating the importance of careful generative model design for estimating behavior policies, and an intuitive notion of complexity for offline RL problems. With its simple interpretation and fewer moving parts, such as no explicit function approximator representing the policy, EMaQ serves as a strong yet easy to implement baseline for future work.

1 INTRODUCTION

Leveraging past interactions in order to improve a decision-making process is the hallmark goal of off-policy reinforcement learning (RL) (Precup et al., 2001; Degris et al., 2012). Effectively learning from past experiences can significantly reduce the amount of online interaction required to learn a good policy, and is a particularly crucial ingredient in settings where interactions are costly or safety is of great importance, such as robotics Gu et al. (2017); Kalashnikov et al. (2018a), health Murphy et al. (2001), dialog agents (Jaques et al., 2019), and education Mandel et al. (2014). In recent years, with neural networks taking a more central role in the RL literature, there have been significant advances in developing off-policy RL algorithms for the function approximator setting, where policies and value functions are represented by neural networks (Mnih et al., 2015; Lillicrap et al., 2015; Gu et al., 2016b;a; Haarnoja et al., 2018; Fujimoto et al., 2018b). Such algorithms, while off-policy in nature, are typically trained in an online setting where algorithm updates are interleaved with additional online interactions. However, in purely offline RL settings, where a dataset of interactions are provided ahead of time and no additional interactions are allowed, the performance of these algorithms degrades drastically (Fujimoto et al., 2018a; Jaques et al., 2019).

A number of recent methods have been developed to address this shortcoming of off-policy RL algorithms. A particular class of algorithms for offline RL that have enjoyed recent success are

those based on dynamic programming and value estimation (Fujimoto et al., 2018a; Jaques et al., 2019; Kumar et al., 2019; Wu et al., 2019; Levine et al., 2020). Most proposed algorithms are designed with a key intuition that it is desirable to prevent policies from deviating too much from the provided collection of interactions. By moving far from the actions taken in the offline data, any subsequently learned policies or value functions may not generalize well and lead to the belief that certain actions will lead to better outcomes than they actually would. Furthermore, due to the dynamics of the MDP, taking out-of-distribution actions may lead to states not covered in the offline data, creating a snowball effect (Ross et al., 2011). In order to prevent learned policies from straying from the offline data, various methods have been introduced for regularizing the policy towards a base behavior policy (e.g. through a divergence penalty (Jaques et al., 2019; Wu et al., 2019; Kumar et al., 2019) or clipping actions (Fujimoto et al., 2018a)).

Taking the above intuitions into consideration, in this work we investigate a simplifcation of the BCQ algorithm (Fujimoto et al., 2018a) (a notable prior work in offline RL), which removes a heuristic design choice and has the property that extracted policies remain exactly within the support of a given behavior policy. In contrast to the theoretical considerations in the original work, we derive this simplified algorithm in a theoretical setup that more closely reflects the resulting algorithm. We introduce the Expected-Max Q-Learning (EMaQ) operator, which interpolates between the standard Q-function evaluation and Q-learning backup operators. The EMaQ operator makes explicit the relation between the proposal distribution and number of samples used, and leads to sub-optimality bounds which introduce a novel notion of complexity for offline RL problems. In its practical implementation for the continuous control and function approximator setting, EMaQ has only two standard components (an estimate of the base behavior policy, and Q functions) and does not explicitly represent a policy, requiring fitting one less function approximator than prior approaches (Fujimoto et al., 2018a; Kumar et al., 2019; Wu et al., 2019).

In online RL, EMaQ is competitive with Soft Actor Critic (SAC) (Haarnoja et al., 2018) and surpasses SAC in the deployment-efficient setting (Matsushima et al., 2020). In the offline RL setting – the main focus of this work – EMaQ matches and outperforms prior state-of-the-art in the D4RL (Fu et al., 2020a) benchmark tasks. Through our explorations with EMaQ we make two intriguing findings. First, due to the strong dependence of EMaQ on the quality of behavior policy used, our results demonstrate the significant impact of careful considerations in modeling the behavior policy that generate the offline interaction datasets. Second, relating to the introduced notion of complexity, in a diverse array of benchmark settings considered in this work we observe that surprisingly little modification to a base behavior policy is necessary to obtain a performant policy. As an example, we observe that while a HalfCheetah random policy obtains a return of 0, a policy that at each state uniformly samples only 5 actions and chooses the one with the best value obtains a return of 2000. The simplicity, intuitive interpretation, and strong empirical performance of EMaQ make it a great test-bed for further examination and theoretical analyses, and an easy to implement yet strong baseline for future work in offline RL.

2 BACKGROUND

Throughout this work, we represent Markov Decision Process (MDP) as $M = \langle \mathcal{S}, \mathcal{A}, r, \mathcal{P}, \gamma \rangle$, with state space $\mathcal{S}$, action space $\mathcal{A}$, reward function $r : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$, transition dynamics $\mathcal{P}$, and discount γ. In offline RL, we assume access to a dataset of interactions with the MDP, which we will represent as collection of tuples $D = \{(s, a, s', r, t)\}^N$, where t is an indicator variable that is set to True when s' is a terminal state. We will use μ to represent the behavior policy used to collect D, and depending on the context, we will overload this notation and use μ to represent an estimate of the true behavior policy. For a given policy π, we will use the notation $d^\pi(s), d^\pi(s, a)$ to represent the state-visitation and state-action visitation distributions respectively.

As alluded to above, a significant challenge of offline RL methods is the problem of distribution shift. At training-time, there is no distribution shift in states as a fixed dataset D is used for training, and the policy and value functions are never evaluated on states outside of $d^\mu(s)$. However, a very significant challenge is the problem of **distribution shift in actions**. Consider the Bellman backup for obtaining the Q-function of a given policy π,

$$\mathcal{T}_\pi Q(s, a) := r(s, a) + \gamma \cdot \mathbb{E}_{s'} \mathbb{E}_{a' \sim \pi(a'|s')} \Big[Q(s', a') \Big] \tag{1}$$

The target Q-values on the right hand side depend on action samples $a' \sim \pi(a'|s')$. If the sampled actions are outside the distribution of actions observed in D, the estimated Q-values can be erroneous leading to incorrect target values. The effects of action distribution shift are further exacerbated in actor-critic algorithms; out of distribution (OOD) actions may incorrectly be assigned high values, in which case the policy will be updated to further sample OOD actions, leading to a hazardous loop.

An important approach – with particular recent interest – to mitigate the effects of both kinds of distributional shift is to devise methods for constraining learned policies to remain close to the behavior policy μ: $d^\pi(s,a) \approx d^\mu(s,a)$. Below, we set the stage by reviewing a closely related prior work in offline RL.

Batch Constrained Q-Learning (BCQ) In BCQ (Fujimoto et al., 2018a) the aim is to constrain a Q-Learning based algorithm such that it will be effective in the offline RL continuous control setting with function approximators. To do so, the trained policy is parameterized as:

$$\pi_\theta(a|s) = \underset{a_i + \xi_\theta(s,a_i)}{\arg\max}\ Q_\psi(s, a_i + \xi_\theta(s,a_i)) \qquad \text{for} \qquad a_i \sim \mu(a|s), i = 1, ..., N \tag{2}$$

$$y(s,a,s',r,t) = \left(r + (1-t)\cdot\gamma \max_{a'_i} Q_{\psi'}(s', a'_i) \right) \qquad \text{for} \qquad a'_i \sim \pi_\theta(a'|s'), i = 1, ..., N \tag{3}$$

$$\mathcal{L}_Q = \left(y(s,a,s',r,t) - Q_\psi(s,a)\right)^2 \tag{4}$$

where $y(s,a)$ are target Q-values, Q_ψ is learned with the objective in equation 4, $\mu(a|s)$ is an estimate of the base behavior policy (a generative model trained using the dataset D), and ξ_θ is an action perturbation model trained to modify actions towards more optimal ones. *Crucially*, each component of the output of ξ_θ is bounded to the range $[-\Phi, \Phi]$. ***The key intuition*** is that because a_i are sampled from an estimate of the behavior policy, they should hopefully be within the distribution observed in D. Thus, since the perturbation model is constrained by the hyperparameter Φ, the perturbed actions should not be too far from actions in the dataset. This should mitigate errors in value estimates, which should in turn lead to better updates for the perturbation model.

3 Expected-Max Q-Learning

We make the observation that, in the BCQ algorithm, if we could obtain a good estimate $\mu(a|s)$ and sufficiently increased the number of samples N, there would be no need for the perturbation network ξ_θ. This simplification would remove one additional function approximator and the associated hyperparameter Φ. This is the driving intuition of our work, which we frame theoretically in a manner that encapsulates the key components: the behavior policy μ and number of samples N. Below, we introduce the Expected-Max Q operator, illustrate its key properties for tabular MDPs, and obtain sub-optimality bounds which can serve as a novel measure of complexity of an offline RL problem for future theoretical work. We then provide an extension to the offline RL setting with function approximators, and then discuss the generative model used to approximate the behavior policy.

3.1 Expected-Max Q Operator

Let $\mu(a|s)$ be an arbitrary behavior policy, and let $\{a_i\}^N \sim \mu(a|s)$ denote sampling N iid actions from $\mu(a|s)$. Let $Q : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$ be an arbitrary function. For a given choice of N, we define the Expected-Max Q-Learning operator (EMaQ) $\mathcal{T}^N_\mu Q$ as follows:

$$\mathcal{T}^N_\mu Q(s,a) := r(s,a) + \gamma \cdot \mathbb{E}_{s'} \mathbb{E}_{\{a'_i\}^N \sim \mu(\cdot|s')} \left[\max_{a' \in \{a'_i\}^N} Q(s', a') \right] \tag{5}$$

$$\text{Q-Evaluation for } \mu \qquad \mathcal{T}_\mu Q(s,a) := r(s,a) + \gamma \cdot \mathbb{E}_{s'} \mathbb{E}_{a' \sim \mu(\cdot|s')} \left[Q(s', a') \right] \tag{6}$$

$$\text{Q-Learning} \qquad \mathcal{T}^* Q(s,a) := r(s,a) + \gamma \cdot \mathbb{E}_{s'} \left[\max_{a'} Q(s', a') \right] \tag{7}$$

This operator provides a natural interpolant between the on-policy backup for μ (Eq. 6) when $N = 1$, and the Q-learning backup (Eq. 7) as $N \to \infty$ (if $\mu(a|s)$ has full support over $\mathcal{A}$). We formalize these observations more precisely below when we articulate the key properties in the tabular MDP setting. We discuss how this relates to existing modified backup operators in the related work.

3.2 DYNAMIC PROGRAMMING PROPERTIES IN THE TABULAR MDP SETTING

To understand any novel backup operator it is useful to first characterize its key dynamic programming properties in the tabular MDP setting. First, we establish that EMaQ retains essential contraction and fixed-point existence properties, regardless of the choice of $N \in \mathbb{N}$. In the interest of space, all missing proofs can be found in Appendix B.

Theorem 3.1. *In the tabular setting, for any $N \in \mathbb{N}$, $\mathcal{T}_\mu^N$ is a contraction operator in the $\mathcal{L}_\infty$ norm. Hence, with repeated applications of the $\mathcal{T}_\mu^N$, any initial Q function converges to a unique fixed point.*

Theorem 3.2. *Let Q_μ^N denote the unique fixed point achieved in Theorem 3.1, and let $\pi_\mu^N(a|s)$ denote the policy that samples N actions from $\mu(a|s)$, $\{a_i\}^N$, and chooses the action with the maximum Q_μ^N. Then Q_μ^N is the Q-value function corresponding to $\pi_\mu^N(a|s)$.*

Proof. (Theorem 3.2) Rearranging the terms in equation 5 we have,

$$\mathcal{T}_\mu^N Q_\mu^N(s,a) = r(s,a) + \gamma \cdot \mathbb{E}_{s'} \mathbb{E}_{a' \sim \pi_\mu^N(a'|s')}[Q_\mu^N(s',a')]$$

Since by definition Q_μ^N is the unique fixed point of $\mathcal{T}_\mu^N$, we have our result. □

From these results we can then rigorously establish the interpolation properties of the EMaQ family.

Theorem 3.3. *Let π_μ^* denote the optimal policy from the class of policies whose actions are restricted to lie within the support of the policy $\mu(a|s)$. Let Q_μ^* denote the Q-value function corresponding to π_μ^*. Furthermore, let Q_μ denote the Q-value function of the policy $\mu(a|s)$. Let $\mu^*(s) := \int_{\textit{Support}(\pi_\mu^*(a|s))} \mu(a|s)$ denote the probability of optimal actions under $\mu(a|s)$. Under the assumption that $\inf_s \mu^*(s) > 0$ and $r(s,a)$ is bounded, we have that,*

$$Q_\mu^1 = Q_\mu \qquad \textit{and} \qquad \lim_{N\to\infty} Q_\mu^N = Q_\mu^*$$

That is, Theorem 3.3 shows that, given a base behavior policy $\mu(a|s)$, the choice of N makes the EMaQ operator interpolate between evaluating the Q-value of μ on the one hand, and learning the optimal Q-value function on the other (optimal subject to the support constraint discussed in Theorem 3.3). In the special case where $\mu(a|s)$ has full support over the action space $\mathcal{A}$, EMaQ interpolates between the standard Q-Evaluation and Q-Learning operators in reinforcement learning.

Intuitively, as we increase N, the fixed-points Q_μ^N should correspond to increasingly better policies $\pi_\mu^N(a|s)$. We show that this is indeed the case.

Theorem 3.4. *For all $N, M \in \mathbb{N}$, where $N > M$, we have that $\forall s \in \mathcal{S}, \forall a \in$ Support$(\mu(\cdot|s))$, $Q_\mu^N(s,a) \geq Q_\mu^M(s,a)$. Hence, $\pi_\mu^N(a|s)$ is at least as good of a policy as $\pi_\mu^M(a|s)$.*

It is also valuable to obtain a sense of how suboptimal $\pi_\mu^N(a|s)$ may be with respect to the optimal policy supported by the policy $\mu(a|s)$.

Theorem 3.5. *For $s \in \mathcal{S}$ let,*

$$\Delta(s,N) = \max_{a \in \text{Support}(\mu(\cdot|s))} Q_\mu^*(s,a) - \mathbb{E}_{\{a_i\}^N \sim \mu(\cdot|s)}[\max_{b \in \{a_i\}^N} Q_\mu^*(s,b)]$$

The suboptimality of Q_μ^N can be upperbounded as follows,

$$\left\|Q_\mu^N - Q_\mu^*\right\|_\infty \leq \frac{\gamma}{1-\gamma} \max_{s,a} \mathbb{E}_{s'}\Big[\Delta(s',N)\Big] \leq \frac{\gamma}{1-\gamma} \max_s \Delta(s,N) \tag{8}$$

The same also holds when Q_μ^* is replaced with Q_μ^N in the definition of Δ.

3.3 A MEASURE OF COMPLEXITY FOR OFFLINE RL

The bounds in equation 8 capture the main intuitions about the interplay between $\mu(a|s)$ and the choice of N. If for each state, $\mu(a|s)$ places sufficient mass over the optimal actions, π_μ^N will be

Algorithm 1: Test-Time Policy π_{test}

Function `TestEnsemble(values)`**:**
  return $\lambda \cdot \min(values) + (1-\lambda) \cdot \max(values)$

Function $\pi_{\text{test}}(s)$**:**
  $\{a_i\}^N \sim \mu(a|s)$
  return $\arg\max_{\{a_i\}^N}$ `TestEnsemble`$\big([Q_i(s,a)$ for all $i]\big)$

close to π^*_μ. The two variants of $\Delta(s, N)$, based on Q^*_μ or Q^N_μ, suggest an intriguing notion of difficulty for an offline RL problem. If we could estimate either of these Q-value functions, then for a desired set of states (such as initial states) we could plot $\Delta(s, N)$ as decreasing function of N. The rate at which this function decreases could serve as an intuitive notion of difficulty for a given offline offline RL problem which consists of an MDP and a given behavior policy. While we leave theoretical investigations of this measure for future work, our empirical results in Section 5 demonstrate that the effective value of N may be surprisingly small.

3.4 Offline RL Setting with Function Approximators

Typically, we are not provided with the policies that generated the provided trajectories. Hence, as a first step we fit a generative model $\mu(a|s)$ to the (s, a) pairs in the offline dataset, representing the mixture of policies that generated this data (details below). Having obtained $\mu(a|s)$, we move on to the EMaQ training procedure. Similar to prior works (Fujimoto et al., 2018a; Kumar et al., 2019; Wu et al., 2019), we train K Q functions (represented by MLPs) and make use of an ensembling procedure to combat overestimation bias (Hasselt, 2010; Van Hasselt et al., 2016; Fujimoto et al., 2018b). Letting D represent the offline dataset, the objective for the Q functions takes the following form:

$$\mathcal{L}(\theta_i) = \mathbb{E}_{(s,a,s',r,t)\sim D}\left[\Big(Q_{\theta_i}(s,a) - y(s,a,s',r,t)\Big)^2\right] \tag{9}$$

$$y(s,a,s',r,t) = \left(r + (1-t)\cdot\gamma \max_{a'_i} Q'_{ens}(s',a'_i)\right) \quad \text{for} \quad a'_i \sim \mu(a'|s'), i = 1, ..., N \tag{10}$$

where t is the indicator variable $\mathbb{1}[s' \text{ is terminal}]$, and Q'_{ens} represents the ensemble of target Q functions. In short, we sample N actions from $\mu(a'|s')$ and take the value of the best action to form the target. The algorithm box describing the full training loop can be viewed in Algorithm 2.

Notably, we do not train an explicit neural network representing the policy. At test-time, given a state s, we sample N actions from $\mu(a|s)$ and choose the action with the maximum value under the ensemble of Q functions (see Algorithm 1). While `TestEnsemble` can differ from the `Ensemble` function used to compute target Q values[1], in this work we used the same ensembling procedure with $\lambda = 1.0$ (with the exception of experiments in Section G).

3.5 Intuition for Offline EMaQ and Modeling $\mu(a|s)$

The intuition for how offline EMaQ aims to address the problem of erroneous value estimates can be understood from attending to equation 10 and the form of the test-time policy (Algorithm 1). In equation 10 we observe that target values for the Q functions are computed by sampling actions **from the behavior policy estimate μ**, and not a separately learned policy that may sample out of distribution actions, as in BCQ. At test-time, our implicit policy is also formed by choosing amongst actions sampled from μ. Hence, if μ accurately estimates the behavior policy well, we will never sample actions outside the support and will not need to evaluate the value of such actions. In the practical setting where μ may have inaccuracies, the hyperparameter N acts as an implicit regularizer: Using very large values of N maximizes the chance of sampling actions that are out

[1] some examples of alternative choices are `mean`, `max`, UCB-style estimates, or simply using just one of the trained Q functions

of distribution and have erroneous value estimates, while smaller N reduces the chance of this happening in every update iteration and therefore smoothens the incorrect values.

With the importance of a good behavior estimates accentuated in our proposed method, we pay closer attention to the choice of generative model used for representing μ. Past works (Fujimoto et al., 2018a; Kumar et al., 2019; Wu et al., 2019) have typically used Variational Auto-Encoders (VAEs) (Kingma & Welling, 2013; Rezende et al., 2014) to represent the behavior distribution $\mu(a|s)$. Unfortunately, after training the aggregate posterior $q_{\text{agg}}(z) := \mathbb{E}_x[q(z|x)]$ of a VAE does not typically align well with its prior, making it challenging to sample from in a manner that effectively covers the distribution it was trained on[2]. We opt for using an autoregressive architecture based on MADE (Germain et al., 2015) as it allows for representing more expressive distributions and enables more accurate sampling. Inspired by recent works (Metz et al., 2017; Van de Wiele et al., 2020), our generative model architecture also makes use of discretization in each action dimension. Full details can be found in Appendix C.

4 Related Work

Relation between EMaQ and BCQ and other sample-max Q-learning methods As noted earlier, the offline EMaQ algorithm can be recovered by removing the perturbation network in BCQ. In constrast to the original work, this simplified form allows us to incorporate the sample-max procedure into our theoretical considerations and leads to a novel notion of complexity for offline RL algorithms. This results in a closer match between the practical algorithm and theory. From the empirical perspective, as we will observe in the experiments below, without our more careful consideration of the often overlooked choice of generative model, this simplified variant of BCQ does not result in a performant policy. Importantly, our theoretical analysis on the family of TD operators described by EMaQ apply beyond BCQ and can also provide new perspectives on some of the highly successful sample-max Q-learning algorithms (Kalashnikov et al., 2018a; Van de Wiele et al., 2020) – particularly on how the proposal distribution affects convergence. **In the interest of space, we have moved the remainder of our "Related Work" section to Appendix A. If our work is accepted in ICLR 2021, we plan to make use of the provided extra page to include this section in the main text of the camera-ready version.**

5 Experiments

For all experiments we make use of the codebase of (Wu et al., 2019), which presents the BRAC off-policy algorithm and examines the importance of various factors in BCQ (Fujimoto et al., 2018a) and BEAR (Kumar et al., 2019) methods. We implement EMaQ into this codebase. We make use of the recently proposed D4RL (Fu et al., 2020b) datasets for bechmarking offline RL.

Online EMaQ Despite obtaining **strong online RL results competitive with and outperforming SAC** (Haarnoja et al., 2018) (Figure 3), as the main focus of our work is for the offline setting, we have placed our online RL methodology and results in Appendix G. However, we emphasize that significance of obtaining strong online RL performance with effectively the same algorithm as the offline setting should not be overlooked. Most prior offline RL works have not considered how their methods might transfer to online or batched online setting, and recent work (Nair et al., 2020) has demonstrated the challenges of finetuning from a policy trained offline, in the online setting.

5.1 Practical Effect of N and the Choice of Generative Model μ

We begin by empirically evaluating key aspects of offline EMaQ, namely the effect of N, and choice of generative model used for representing the behavior estimate μ. In prior approaches such as those described in the background section of this work, care must be taken in choosing the hyperparameter that dictates the extent to which learned policies can deviate from the base behavior policies; too small and we cannot improve upon the base policy, too large and the value of actions cannot be correctly estimated. In EMaQ, at least in theory, choosing higher values of N should result in

[2]past works typically clip the range of the latent variable z and adjust the weighting of the KL term in the evidence lower-bound to ameliorate the situation

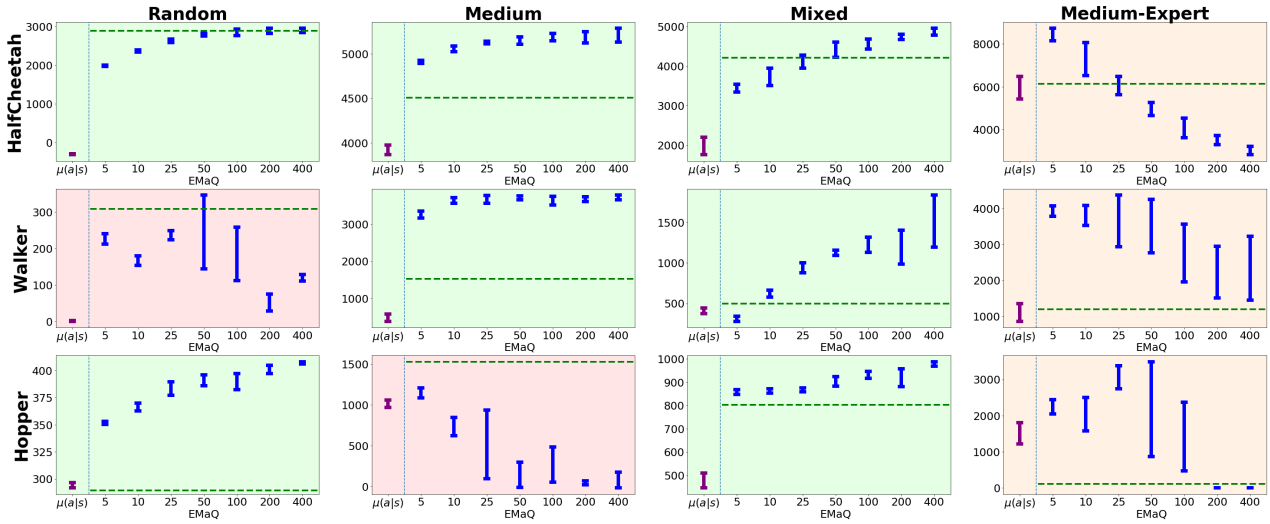

Figure 1: Results for evaluating EMaQ on D4RL benchmark's (Fu et al., 2020b) standard Mujoco domains, with $N \in \{5, 10, 25, 50, 100, 200, 400\}$. Values above $\mu(a|s)$ represent the result of evaluating the base behavior policies. Horizontal green lines represent the reported performance of BEAR in the D4RL benchmark (apples to apples comparisons in Figure 2). The types of offline datasets are: **random**: 1M transitions are collected by a random agent, **medium**: 1M transitions are collected by a half-trained SAC (Haarnoja et al., 2018) policy, **mixed**: the replay buffer of this half-trained policy, and **medium-expert**: combination of the medium dataset and 1M additional transitions from a fully trained policy. Refer to main text (Section 5.1) for description of color-coding. For better legibility, we have included a larger variant of these plots in the Appendix K. Full experimental details in Appendix H.

strictly better policies. Additionally, there exists a concern that N may need to be impractically large. Thus, we empirically investigate to what extent the monotonic trend holds in practice, and seek to understand what magnitudes of N result in good policies in practical benchmark domains. Figure 1 presents our results with $N \in \{5, 10, 25, 50, 100, 200, 400\}$. In the green plots, we observe that empirical results follow our intuitions: with increasing N the resultant policies become better. In the medium-expert settings (i.e. orange plots), while for smaller values of N we observe strong performance, there appears to be a downward trend. As discussed in Section 3.5, smaller values of N result in an implicit regularization. Hence, the orange plots may indicate that even with the stronger choice of generative models in EMaQ, inaccuracies in value estimates may still exist, suggesting the need for future work that introduces better regularizers for the value functions than ensembling (Kumar et al., 2020). Lastly, the red plots indicate settings where behavior is erratic. Closer examination of training curves and our experiments with other off-policy methods (Figure 2) suggests that this may be due to the intrinsic nature of these environment and data settings.

The dashed horizontal lines in Figure 2 represent the performance of BEAR – which uses a VAE for representing μ – as reported in the D4RL (Fu et al., 2020b) benchmark paper (apples to apples comparison in Section 5.2). Our results demonstrate that the combination of a strong generative model and EMaQ's simply constrained backup operator can match and in many cases noticeably exceed results from prior algorithms and design choices. Comparing Figure 2 to Figure 6 in the Appendix, **we observe that our choice of generative model is crucial to the performance of EMaQ**. With a VAE architecture as used in prior work, EMaQ's performance is significantly reduced, in most cases worse than prior reported results for BEAR, and never exhibits a monotonic trend as a function of N. This is despite the fact that when evaluating the performance of the behavior estimate μ under the two architecture choices results in almost identical results (the first column of each sub-plot corresponding to $\mu(a|s)$). We do not believe that autoregressive models are intrinsically better than VAEs, but rather our results demonstrate the need for more careful attention on the choice of $\mu(a|s)$. **Since EMaQ is closely tied to the choice of behavior model, it may be more effective for evaluating how well $\mu(a|s)$ represents the given offline dataset. From a practical perspective, our results suggest that for a given domain, focusing efforts on building-in good inductive biases in the generative models and value functions might be sufficient to obtain strong offline RL performance in many domains.**

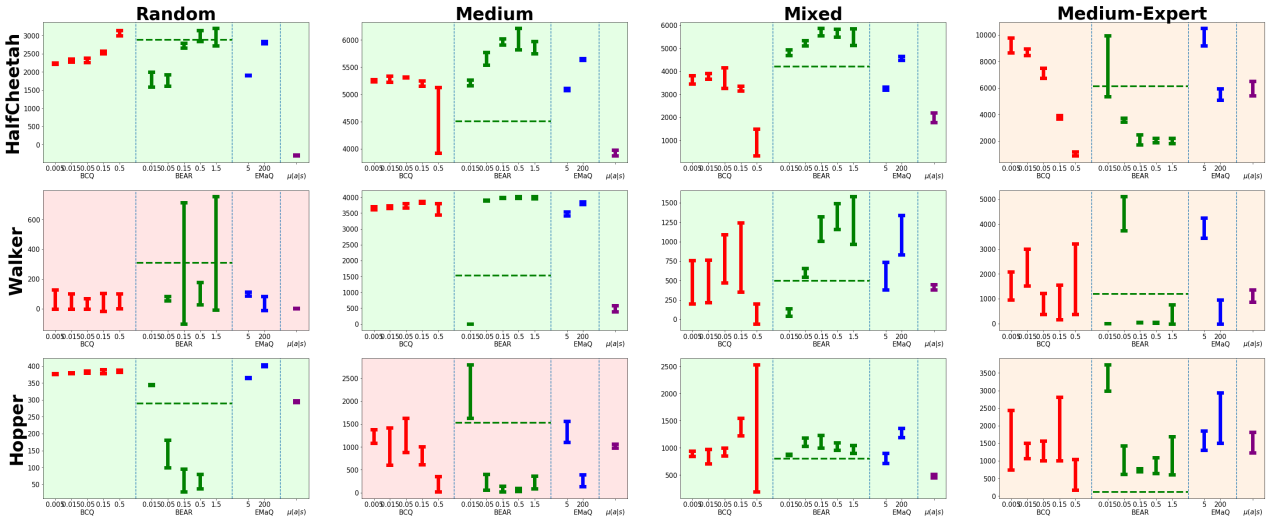

Figure 2: Comparison of EMaQ, BCQ, and BEAR on D4RL (Fu et al., 2020b) benchmark domains when using our proposed autoregressive $\mu(a|s)$. For both BCQ and BEAR, from left to right as the value of the hyperparameter increases, the allowed deviation from $\mu(a|s)$ increases. Horizontal green lines represent the reported performance of BEAR in the D4RL benchmark. Color-coding follows Figure 1. For better legibility, we have included a larger variant of these plots in the Appendix K. Full experimental details in Appendix H.

5.2 Comparison on D4RL Offline RL Benchmark

To evaluate EMaQ with respect to prior methods, we compare to two popular and closely related prior methods for offline RL, BCQ (Fujimoto et al., 2018a) and BEAR (Kumar et al., 2019). As with the previous section, full experimental details can be found in Appendix H.1. Figure 2 and Table 3 (Appendix H.3) present our empirical results. Note that with our proposed autoregressive models, the results for BEAR are matched and in some cases noticeably above the values reported in the D4RL benchmark (Fu et al., 2020b) (green horizontal lines). For easier interpretation, the plots are colored the same as in Figure 1. Our key take-away is that despite its simplistic form, EMaQ is strongly competitive with prior state-of-the-art methods, and in the case of Table 3 (Appendix H.3) outperforms prior approaches. Despite this, there remain many domains in the D4RL benchmark on which none of the considered algorithms make any progress (Table 3, Appendix H.3), indicating that much algorithmic advances are still necessary for solving many of the considered domains.

A very eye-catching result in above figures is that in almost all settings of the standard Mujoco environments (Figures 1 and 2), just $N = 5$ significantly improves upon $\mu(a|s)$ and in most settings matches or exceeds significantly beyond previously reported results. Concretely, this means that in the HalfCheetah-Random setting, if at each state we sample 5 actions uniformly random and choose the best one under the learned Q-value function, we convert a random policy with return 0 to a policy with return 2000. **In this way, EMaQ provides a quite intuitive and surprising measure of the complexity for offline RL problems. This empirical observation also corroborates our discussion in Section 3.3, encouraging future theoretical investigations into $\boldsymbol{\Delta(s, N)}$.**

6 Conclusion

In this work, we investigate a significant simplification of the BCQ (Fujimoto et al., 2018a) algorithm by removing the heuristic perturbation network. By introducing the Expect-Max Q-Learning operator, we present a novel theoretical setup that takes into account the proposal distribution $\mu(a|s)$ and the number of action samples N, and hence more closely matches the resulting practical algorithm. With fewer moving parts and one less function approximator, EMaQ matches and outperforms prior state-of-the-art in online and offline RL. Our investigations with EMaQ demonstrate the significance of careful considerations in the design of generative models used. Furthermore, our theoretical and empirical findings bring into light novel notions of complexity for offline RL problems. Given the simplicity, tractable theory, and state-of-the-art performance of EMaQ, we hope our work can serve as a foundation for future works on understanding and improving offline RL.

REFERENCES

Brandon Amos, Lei Xu, and J Zico Kolter. Input convex neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 146–155. JMLR. org, 2017.

Marc G Bellemare, Georg Ostrovski, Arthur Guez, Philip S Thomas, and Rémi Munos. Increasing the action gap: New operators for reinforcement learning. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.

Richard Y Chen, Szymon Sidor, Pieter Abbeel, and John Schulman. Ucb exploration via q-ensembles. *arXiv preprint arXiv:1706.01502*, 2017.

Thomas Degris, Martha White, and Richard S Sutton. Off-policy actor-critic. *arXiv preprint arXiv:1205.4839*, 2012.

Amir-massoud Farahmand. Action-gap phenomenon in reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 172–180, 2011.

Roy Fox, Ari Pakman, and Naftali Tishby. Taming the noise in reinforcement learning via soft updates. *arXiv preprint arXiv:1512.08562*, 2015.

Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. D4rl: Datasets for deep data-driven reinforcement learning, 2020a.

Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. Datasets for data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*, 2020b.

Scott Fujimoto, David Meger, and Doina Precup. Off-policy deep reinforcement learning without exploration. *arXiv preprint arXiv:1812.02900*, 2018a.

Scott Fujimoto, Herke Van Hoof, and David Meger. Addressing function approximation error in actor-critic methods. *arXiv preprint arXiv:1802.09477*, 2018b.

Mathieu Germain, Karol Gregor, Iain Murray, and Hugo Larochelle. Made: Masked autoencoder for distribution estimation. In *International Conference on Machine Learning*, pp. 881–889, 2015.

Shixiang Gu, Timothy Lillicrap, Zoubin Ghahramani, Richard E Turner, and Sergey Levine. Q-prop: Sample-efficient policy gradient with an off-policy critic. *arXiv preprint arXiv:1611.02247*, 2016a.

Shixiang Gu, Timothy Lillicrap, Ilya Sutskever, and Sergey Levine. Continuous deep q-learning with model-based acceleration. In *International Conference on Machine Learning*, 2016b.

Shixiang Gu, Ethan Holly, Timothy Lillicrap, and Sergey Levine. Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In *International Conference on Robotics and Automation*, 2017.

Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 2018.

Hado V Hasselt. Double q-learning. In *Advances in neural information processing systems*, pp. 2613–2621, 2010.

Natasha Jaques, Shixiang Gu, Dzmitry Bahdanau, José Miguel Hernández-Lobato, Richard E Turner, and Douglas Eck. Sequence tutor: Conservative fine-tuning of sequence generation models with kl-control. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1645–1654. JMLR. org, 2017.

Natasha Jaques, Asma Ghandeharioun, Judy Hanwen Shen, Craig Ferguson, Agata Lapedriza, Noah Jones, Shixiang Gu, and Rosalind Picard. Way off-policy batch deep reinforcement learning of implicit human preferences in dialog. *arXiv preprint arXiv:1907.00456*, 2019.

Sham M Kakade. A natural policy gradient. In *Advances in neural information processing systems*, pp. 1531–1538, 2002.

Dmitry Kalashnikov, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre Quillen, Ethan Holly, Mrinal Kalakrishnan, Vincent Vanhoucke, and Sergey Levine. Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation. In *Conference on Robot Learning*, 2018a.

Dmitry Kalashnikov, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre Quillen, Ethan Holly, Mrinal Kalakrishnan, Vincent Vanhoucke, et al. Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation. *arXiv preprint arXiv:1806.10293*, 2018b.

Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

Aviral Kumar, Justin Fu, Matthew Soh, George Tucker, and Sergey Levine. Stabilizing off-policy q-learning via bootstrapping error reduction. In *Advances in Neural Information Processing Systems*, 2019.

Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. Conservative q-learning for offline reinforcement learning. *arXiv preprint arXiv:2006.04779*, 2020.

Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.

Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.

Travis Mandel, Yun-En Liu, Sergey Levine, Emma Brunskill, and Zoran Popovic. Offline policy evaluation across representations with applications to educational games. In *International Conference on Autonomous Agents and Multiagent Systems*, 2014.

Tatsuya Matsushima, Hiroki Furuta, Yutaka Matsuo, Ofir Nachum, and Shixiang Gu. Deployment-efficient reinforcement learning via model-based offline optimization. *arXiv preprint arXiv:2006.03647*, 2020.

Luke Metz, Julian Ibarz, Navdeep Jaitly, and James Davidson. Discrete sequential prediction of continuous actions for deep rl. *arXiv preprint arXiv:1705.05035*, 2017.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

Susan A Murphy, Mark J van der Laan, James M Robins, and Conduct Problems Prevention Research Group. Marginal mean models for dynamic regimes. *Journal of the American Statistical Association*, 2001.

Ofir Nachum, Mohammad Norouzi, and Dale Schuurmans. Bridging the gap between value and policy based reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 2772–2782, 2017.

Ashvin Nair, Murtaza Dalal, Abhishek Gupta, and Sergey Levine. Accelerating online reinforcement learning with offline datasets. *arXiv preprint arXiv:2006.09359*, 2020.

Jan Peters, Katharina Mulling, and Yasemin Altun. Relative entropy policy search. In *Twenty-Fourth AAAI Conference on Artificial Intelligence*, 2010.

Doina Precup, Richard S Sutton, and Sanjoy Dasgupta. Off-policy temporal-difference learning with function approximation. In *International Conference on Machine Learning*, 2001.

Konrad Rawlik, Marc Toussaint, and Sethu Vijayakumar. On stochastic optimal control and reinforcement learning by approximate inference. In *Twenty-Third International Joint Conference on Artificial Intelligence*, 2013.

Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.

Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 627–635, 2011.

Reuven Y Rubinstein and Dirk P Kroese. *The cross-entropy method: a unified approach to combinatorial optimization, Monte-Carlo simulation and machine learning*. Springer Science & Business Media, 2013.

Moonkyung Ryu, Yinlam Chow, Ross Anderson, Christian Tjandraatmadja, and Craig Boutilier. Caql: Continuous action q-learning. *arXiv preprint arXiv:1909.12397*, 2019.

John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pp. 1889–1897, 2015.

Tom Van de Wiele, David Warde-Farley, Andriy Mnih, and Volodymyr Mnih. Q-learning in enormous action spaces via amortized approximate maximization. *arXiv preprint arXiv:2001.08116*, 2020.

Hado Van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. In *Thirtieth AAAI conference on artificial intelligence*, 2016.

Lilian Weng. Exploration strategies in deep reinforcement learning, Jun 2020. URL `https://lilianweng.github.io/lil-log/2020/06/07/exploration-strategies-in-deep-reinforcement-learning.html`.

Yifan Wu, George Tucker, and Ofir Nachum. Behavior regularized offline reinforcement learning. *arXiv preprint arXiv:1911.11361*, 2019.