
Accountable Off-Policy Evaluation With Kernel Bellman Statistics

Yihao Feng¹ Tongzheng Ren^{*1} Ziyang Tang^{*1} Qiang Liu¹

Abstract

We consider off-policy evaluation (OPE), which evaluates the performance of a new policy from observed data collected from previous experiments, without requiring the execution of the new policy. This finds important applications in areas with high execution cost or safety concerns, such as medical diagnosis, recommendation systems and robotics. In practice, due to the limited information from off-policy data, it is highly desirable to construct rigorous confidence intervals, not just point estimation, for the policy performance. In this work, we propose a new *variational* framework which reduces the problem of calculating tight confidence bounds in OPE into an optimization problem on a feasible set that catches the true state-action value function with high probability. The feasible set is constructed by leveraging statistical properties of a recently proposed kernel Bellman loss (Feng et al., 2019). We design an efficient computational approach for calculating our bounds, and extend it to perform post-hoc diagnosis and correction for existing estimators. Empirical results show that our method yields tight confidence intervals in different settings.

1. Introduction

Reinforcement learning (Sutton & Barto, 1998) has achieved remarkable successes recently, but is still highly data demanding. An essential approach to improve data efficiency is to use *off-policy* methods, which leverage historical data collected from different behavior policies (a.k.a. off-policy data) to evaluate and optimize the target policies of interest. Using off-policy data is of critical importance in many real-world applications (e.g., Murphy et al., 2001; Li et al., 2011; Bottou et al., 2013; Thomas et al., 2017), especially in cases when it is infeasible or even risky to collect online

experimental data or build high-quality simulators, such as robotics, advertisement, online education, and medical treatment (e.g., Murphy et al., 2001; Li et al., 2011; Hirano et al., 2003; Liao et al., 2019).

This work concerns the problem of off-policy evaluation (OPE), which estimates the expected reward of a given target policy from off-policy data. Because the information in off-policy data is often limited, typical point estimation may suffer from large error in difficult cases. Therefore, for making high-stakes decisions in areas such as medical treatment, it is crucially important to provide reliable *confidence bounds* for quantifying the uncertainties in the estimation.

Importance sampling (IS) provides a basic principle for developing OPE methods (e.g., Dudík et al., 2011; Jiang & Li, 2016; Thomas & Brunskill, 2016; Liu et al., 2018; Nachum et al., 2019; Zhang et al., 2020), as well as doubly robust estimators (Dudík et al., 2011; Jiang & Li, 2016; Farajtabar et al., 2018; Kallus & Uehara, 2020; Tang et al., 2020), as well as constructing confidence bounds for OPE (e.g., Thomas et al., 2015a; Hanna et al., 2017). However, the typical IS-based methods suffer from a number of critical disadvantages:

- *Requiring known behavior policies*: The typical IS estimators require the observed data to be drawn from a known behavior policy. However, this is often not the case in practice. In practice, the policies from which the observed data is collected may be unknown, or even not exist (e.g., when the data is drawn from a mix of different policies).
- *High variance and loose bounds*: The variance of IS-based estimators can be excessively high when the target and behavior policies are very different from each other. As a consequence, concentration confidence bounds based on IS (e.g., Thomas et al., 2015a;b) could be very loose.
- *Curse of long horizon*: The variance of IS estimators are known to deteriorate quickly when the horizon the Markov decision processes becomes large. Although there are recent improved methods for long horizon problems (Liu et al., 2018; Tang et al., 2020; Mousavi et al., 2020), they introduce additional biases due to the estimation of the density ratios.

Our Method We propose a variational framework for constructing *tight, provable confidence bounds* which avoids both the requirement on known behavior policies and the

^{*}Equal contribution ¹Department of Computer Science, The University of Texas at Austin. Correspondence to: Yihao Feng <yihao@cs.utexas.edu>, Tongzheng Ren <tzren@cs.utexas.edu>.

curse of long horizon. Instead of using importance sampling, we construct bounds based on the state-action value function (a.k.a. Q-function) specified by Bellman equation. The idea is to construct a high confidence feasible set of the true Q-function of the target policy, by leveraging the tail bound of the kernel Bellman loss (Feng et al., 2019), and derive the upper (resp. lower) bounds of the expected reward of interest by solving a maximization (resp. minimization) problem on the high confidence feasible set. Further, by constraining the optimization inside a reproducing kernel Hilbert space (RKHS) ball, we obtain a computationally efficient framework for providing confidence bounds of off-policy value evaluation.

Post-hoc Analysis In practice, we may already have an estimated Q-function from previous data/experiments, and we want to construct post-hoc confidence bounds around the existing estimation, or perform post-hoc correction if the estimation is detected to deviate significantly from the true value. In this work, we enable such post-hoc diagnosis by extending our main framework. We empirically evaluate our method on several control benchmarks; results show that our method can provide tight confidence bounds, and also reasonable post-hoc corrections for existing estimators.

2. Related Work

Off-policy value (point) estimation has been studied in various reinforcement learning settings, including contextual bandits (e.g., Dudík et al., 2011; Wang et al., 2017), finite horizon Markov decision processes (MDPs) (e.g., Thomas & Brunskill, 2016; Jiang & Li, 2016; Farajtabar et al., 2018; Xie et al., 2019), and infinite horizon RL MDPs (e.g., Liu et al., 2018; Nachum et al., 2019; Tang et al., 2020). Unlike our work, most of these works are not behavior agnostic and require to know the behavior policy, except a line of very recent works (e.g., Nachum et al., 2019; Zhang et al., 2020; Mousavi et al., 2020).

A smaller set of works focus on confidence bounds for various RL problems (e.g., White & White, 2010; Thomas et al., 2015a;b; Hanna et al., 2017; Dann et al., 2019; Jiang & Huang, 2020; Duan et al., 2020). A recent line of works (e.g., Dann et al., 2019; Jin et al., 2018; 2020) have utilized upper confidence bounds for explorations on tabular (or linear) MDPs. Thomas et al. (2015a); Hanna et al. (2017) give high concentration confidence bounds for off-policy estimators, based on importance sampling methods, which suffer from the issue of high variance of IS-based estimators.

Recently, Jiang & Huang (2020) proposes upper/lower bounds for OPE based a minimax formulation of double robustness (e.g., Tang et al., 2020; Uehara et al., 2020), but their method does not take the randomness of empirical data (data uncertainty) into account to get high confidence

bounds. Concurrently, Duan et al. (2020) provides a data-dependent confidence bounds, based on regression-based Fitted Q iteration (FQI), by assuming the function class is linear and making additional assumptions on MDPs.

Compared with these existing works, our method provides tighter bounds that avoid the high variance issue of IS, work for behavior-agnostic OPE estimators, and can be directly applied to both continuous and discrete state and action MDPs once a kernel function can be properly defined.

Finally, the kernel method has been widely used in reinforcement learning, but is typically used to represent value functions or transition models (e.g., Xu et al., 2005; 2007; Jung & Polani, 2007; Taylor & Parr, 2009; Grunewalder et al., 2012; Farahmand et al., 2016). Our work admits a novel application of kernel in RL, leveraging it as a technical tool for both defining the kernel Bellman loss and constructing high confidence bounds.

3. Off-Policy Evaluation and Kernel Bellman Statistics

We start with introducing background on reinforcement learning and off-policy evaluation (OPE), then give a review on the kernel Bellman loss (Feng et al., 2019) and introduce the concentration bounds for it.

3.1. Preliminary

Denote by $M = \langle \mathcal{S}, \mathcal{A}, \mathcal{P}, r, \gamma, \mu_0 \rangle$ a Markov decision process (MDP), where \mathcal{S} is the state space; \mathcal{A} is the action space; $\mathcal{P}(s'|s, a)$ is the transition probability; $r(s, a)$ is the average immediate reward; $\mu_0 \in \Delta(\mathcal{S})$ is the initial state distribution, and $\gamma \in (0, 1)$ is the discount factor. We focus on the discounted case ($\gamma < 1$) in the main paper, and discuss the extension to average reward case ($\gamma = 1$) in Appendix C.

A policy π specifies a distribution of actions given states, and $\pi(a|s)$ denotes the probability of selecting a given s . Given a *target* policy π , we are interested in estimating the expected total discounted reward associated with π :

$$\eta^\pi := \lim_{T \rightarrow \infty} \mathbb{E}_\pi \left[\sum_{t=0}^T \gamma^t r_t \right] = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t r_t \right], \quad (1)$$

where we start from an initial state $s_0 \sim \mu_0$, and execute the policy π in the MDP to generate trajectories; here T is the trajectory length. We mainly consider the case with infinite horizon ($T = \infty$) in this work.

Assume we are given an observed set of transition pairs $\mathcal{D} = (s_i, a_i, r_i, s'_i)_{i=1}^n$, where r_i and s'_i are the observed local reward and the next state following (s_i, a_i) . The data is assumed to be *off-policy* in that they are collected from some arbitrary, unknown policy (a.k.a. behavior policy),

which is different from the target policy π of interest. The goal of off-policy evaluation (OPE) is to provide interval estimation (upper and lower bounds) for the expected reward η^π from the off-policy data \mathcal{D} , without knowing the underlying behavior policy (i.e., behavior agnostic).

Value Estimation via Q-Function We approach this problem by approximating the expected reward η^π through the Q-function. The Q-function (a.k.a state-action value function) $Q^\pi: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ of a policy π is defined by

$$Q^\pi(s, a) := \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \mid s_0 = s, a_0 = a \right],$$

which equals the expected long-term discounted reward when we execute policy π starting from (s, a) .

Obviously, Q^π and η^π are related via

$$\eta^\pi = \eta(Q^\pi) := \mathbb{E}_{s_0 \sim \mu_0, a_0 \sim \pi(\cdot | s_0)} [Q^\pi(s_0, a_0)], \quad (2)$$

where s_0 is drawn from the fixed initial distribution μ_0 . Therefore, given an empirical approximation \hat{Q} of Q^π , we can estimate η^π by $\eta(\hat{Q})$, which can be further approximated by drawing an i.i.d. sample $(s_{i,0}, a_{i,0})_{i=1}^N$ from $\mu_0 \times \pi$:

$$\eta(\hat{Q}) \approx \hat{\eta}(\hat{Q}) := \frac{1}{N} \sum_{i=1}^N \hat{Q}(s_{i,0}, a_{i,0}). \quad (3)$$

Since μ_0 and π are assumed to be known, we can draw an i.i.d. sample with a very large size N to get arbitrarily high accuracy in the Monte Carlo estimation.

Bellman Equation It is well-known that $Q = Q^\pi$ is the unique solution of the *Bellman equation* (Puterman, 1994): $Q = \mathcal{B}_\pi Q$, where \mathcal{B}_π is the *Bellman evaluation operator*, defined by

$$\mathcal{B}_\pi Q(s, a) := \mathbb{E}_{(s', a')} [r(s, a) + \gamma Q(s', a') \mid s, a].$$

where (s', a') is drawn from $s' \sim \mathcal{P}(\cdot | s, a)$ and $a' \sim \pi(\cdot | s')$. For simplicity, we can define the *Bellman residual operator* as $\mathcal{R}_\pi Q = \mathcal{B}_\pi Q - Q$.

We can approximate the Bellman operator on a state-action pair (s_i, a_i) by bootstrapping:

$$\hat{\mathcal{B}}_\pi Q(s_i, a_i) := r_i + \gamma \mathbb{E}_{a' \sim \pi(\cdot | s'_i)} [Q(s'_i, a')],$$

where r_i, s'_i are the observed local reward and the next state of (s_i, a_i) , and a' is drawn from policy π by Monte Carlo (same as expected Sarsa (Sutton & Barto, 1998)). Because the policy is known, we can approximate the expectation term $\mathbb{E}_{a' \sim \pi(\cdot | s'_i)} [Q(s'_i, a')]$ upto arbitrarily high accuracy by drawing a large sample from $\pi(\cdot | s'_i)$. We can also empirically estimate the Bellman residual by

$\hat{\mathcal{R}}_\pi Q(s_i, a_i) := \hat{\mathcal{B}}_\pi Q(s_i, a_i) - Q(s_i, a_i)$, which provides an unbiased estimation of the Bellman residual, that is, $\mathbb{E}[\hat{\mathcal{R}}_\pi Q(s_i, a_i) \mid s_i, a_i] = \mathcal{R}_\pi Q(s_i, a_i)$. Note that $\hat{\mathcal{B}}_\pi Q(s_i, a_i)$ is a function of (s_i, a_i, r_i, s'_i) , but the dependency on r_i, s'_i is dropped in notation for convenience.

The Double Sampling Problem A naive approach for estimating Q^π is to minimize the empirical Bellman residual. However, a key barrier is the so-called ‘‘double sampling’’ problem. To illustrate, consider estimating Q^π by minimizing the square norm of the empirical Bellman residual:

$$\hat{L}_2(Q) = \frac{1}{n} \sum_{i=1}^n \left[\left(\hat{\mathcal{R}}_\pi Q(s_i, a_i) \right)^2 \right].$$

Unfortunately, minimizing $\hat{L}_2(Q)$ does not yield a consistent estimation of Q^π , because $(\hat{\mathcal{R}}_\pi Q(s_i, a_i))^2$ does not correctly estimate $(\mathcal{R}_\pi Q(s_i, a_i))^2$ due to the square function outside of the empirical Bellman residual (even though $\hat{\mathcal{R}}_\pi Q(s_i, a_i)$ is an unbiased estimation of $\mathcal{R}_\pi Q(s_i, a_i)$). One strategy for obtaining an unbiased estimation of the exact square $(\mathcal{R}_\pi Q(s_i, a_i))^2$ is to multiply two copies of $\hat{\mathcal{R}}_\pi Q(s_i, a_i)$ obtained from two *independent* copies of r_i, s'_i following the same (s_i, a_i) , hence requiring *double sampling* (Baird, 1995). However, double sampling rarely happens in observed data collected from natural environment, and is impractical to use in real-world problems. As a result, many works on Q-function estimation resort to temporal difference (TD) based algorithms, which, however, suffer from instability and divergence issues, especially when nonlinear function approximation is used.

3.2. Kernel Bellman Statistics

Recently, Feng et al. (2019) proposed a kernel Bellman loss to address the double sampling problem in value function estimation. This objective function provides a consistent estimation of the derivation from a given Q to the true Q^π , and hence allows us to estimate Q^π without requiring double sampling, nor risking the instability or divergence in TD-based algorithms. We now give a brief introduction to the kernel Bellman loss and the related U/V-statistics, and their concentration inequalities that are useful in our framework.

For notation, we use $x = (s, a)$ to denote a state-action pair, and $x' := (s', a')$ a successive state-action pair following $x = (s, a)$ collected under some behavior policy. We use $\tau = (s, a, r, s')$ to denote a transition pair.

Following Feng et al. (2019), let $K: (\mathcal{S} \times \mathcal{A}) \times (\mathcal{S} \times \mathcal{A}) \rightarrow \mathbb{R}$ be an *integrally strictly positive definite (ISPD) kernel*. The expected kernel Bellman loss is defined by

$$L_K(Q) := \mathbb{E}_{x, \bar{x} \sim \mu} [\mathcal{R}_\pi Q(x) K(x, \bar{x}) \mathcal{R}_\pi Q(\bar{x})], \quad (4)$$

where $\bar{x} = (\bar{s}, \bar{a})$ is an independent and identical copy of x , and μ is any distribution supported on $\mathcal{S} \times \mathcal{A}$. Note that μ

can be either the on-policy visitation distribution induced by π , or some other valid distributions defined by the observed data (the *off-policy* case).

As shown in Feng et al. (2019), the expected kernel Bellman loss fully characterizes the true Q^π function in that $L_K(Q) \geq 0$ for any Q , and

$$L_K(Q) = 0 \quad \text{if and only if} \quad Q = Q^\pi.$$

Another key property of $L_K(Q)$ is that it can be easily estimated and optimized from observed transitions, without requiring double samples. Given a set of observed transition pairs $\mathcal{D} = \{\tau_i\}_{i=1}^n$, we can use the so-called *V-statistics* to estimate $L_K(Q)$:

$$\widehat{L}_K^V(Q) := \frac{1}{n^2} \sum_{i,j=1}^n \ell_{\pi,Q}(\tau_i, \tau_j), \quad (5)$$

where $\ell_{\pi,Q}$ is defined by

$$\ell_{\pi,Q}(\tau_i, \tau_j) := \widehat{\mathcal{R}}_\pi Q(x_i) K(x_i, x_j) \widehat{\mathcal{R}}_\pi Q(x_j). \quad (6)$$

Another way is to use the so-called *U-statistics*, which removes the diagonal terms in (5):

$$\widehat{L}_K^U(Q) := \frac{1}{n(n-1)} \sum_{1 \leq i \neq j \leq n} \ell_{\pi,Q}(\tau_i, \tau_j). \quad (7)$$

We call $\widehat{L}_K^U(Q)$ and $\widehat{L}_K^V(Q)$ the kernel Bellman U/V-statistics. Both the U/V-statistics can be shown to give consistent estimation of $L_K(Q)$. The U-statistics $\widehat{L}_K^U(Q)$ is known to be an unbiased estimator of the kernel loss $L_K(Q)$, but can be negative and hence unstable when being used as an objective function in practice. In comparison, the V-statistics $\widehat{L}_K^V(Q)$ is always non-negative, that is, $\widehat{L}_K^V(Q) \geq 0$ for any Q , and hence behaves more stably in practice. We mainly use the V-statistics in this work.

3.3. Concentration Bounds of Kernel Bellman Statistics

We now introduce concentration inequalities for the kernel Bellman statistics, which form an important tool in our framework. The inequalities here are classical results from Hoeffding (1963).

Proposition 3.1 (Concentration bound of U-/V-statistics). *Consider a set of i.i.d. random transition pairs $\{\tau_i\}_{i=1}^n$ with $x_i \sim \mu$. Assume*

$$\sup_{\tau, \bar{\tau}} |\ell_{\pi,Q}(\tau, \bar{\tau})| \leq \ell_{\max} < \infty,$$

and n is an even number, then we have the following concentration bound for U-statistics,

$$\mathbb{P} \left[\left| \widehat{L}_K^U(Q) - L_K(Q) \right| \geq 2\ell_{\max} \sqrt{\frac{\log \frac{2}{\delta}}{n}} \right] \leq \delta,$$

and for V-statistics,

$$\mathbb{P} \left[\left| \widehat{L}_K^V(Q) - L_K(Q) \right| \geq 2\ell_{\max} \left(\frac{n-1}{n} \sqrt{\frac{\log \frac{2}{\delta}}{n}} + \frac{1}{n} \right) \right] \leq \delta.$$

We include the proof in the Appendix A for completeness. In addition, in the case when $Q = Q^\pi$, it can be shown that the i.i.d. assumption can be removed; see Appendix A.3.

The key quantity in the concentration inequalities is the upper bound ℓ_{\max} . The lemma below shows that it can be estimated easily in practice.

Lemma 3.1. *Assume the reward function and the kernel function are bounded, i.e. $\sup_x |r(x)| \leq r_{\max}$, $\sup_{x, \bar{x}} |K(x, \bar{x})| \leq K_{\max}$. Then we have*

$$\begin{aligned} \sup_x |Q^\pi(x)| &\leq Q_{\max} := \frac{r_{\max}}{1-\gamma}, \\ \sup_{\tau, \bar{\tau}} |\ell_{\pi, Q^\pi}(\tau, \bar{\tau})| &\leq \ell_{\max} := \frac{4K_{\max}r_{\max}^2}{(1-\gamma)^2}. \end{aligned}$$

4. Accountable Off-Policy Evaluation

This section introduces our approach for accountable off-policy evaluation. We start with the main variational framework for providing confidence bounds in Section 4.1, and then discuss post-hoc diagnosis analysis and correction for existing estimators in Section 4.3 and 4.4.

4.1. Variational Confidence Bounds for Off-Policy Evaluation

We consider the problem of providing confidence bounds for the expected reward η^π . To simplify the presentation, we focus on the upper bounds here, and one can derive the lower bounds with almost the same arguments.

Let \mathcal{F} be a function class that contains Q^π , that is, $Q^\pi \in \mathcal{F}$. Then we can construct an upper bound of η^π by solving the following variational (or functional) optimization on \mathcal{F} :

$$\eta^+ = \max_{Q \in \mathcal{F}} \left\{ \eta(Q), \text{ s.t. } L_K(Q) \leq \lambda \right\}, \quad \forall \lambda > 0. \quad (8)$$

Here η^+ is an upper bound of η^π because Q^π satisfies the condition that $Q^\pi \in \mathcal{F}$ and $L_K(Q^\pi) = 0$, and hence $\eta^+ \geq \eta^\pi$ follows the definition of the max operator.

In practice, we can not exactly calculate η^+ , as it involves the exact kernel loss $L_K(Q)$ and the exact $\eta(Q)$. However, if we replace $L_K(Q)$ and $\eta(Q)$ with proper empirical estimates, we can derive a computationally tractable high-probability upper bound of η^π .

Algorithm 1 Confidence Bounds for Off-Policy Evaluation

Input: Off-policy data $\mathcal{D} = \{x_i, r_i, s'_i\}_{i \leq n}$; maximum reward r_{\max} , discounted factor γ , positive definite kernel K , RKHS norm ρ , random feature $\Phi(\cdot)$, failure probability δ , Monte Carlo sample size N .

Draw $\{x_{i,0}\}_{i=1}^N$ from the initial distribution $\mu_0 \times \pi$; decide λ_K and λ_η by concentration inequalities.

Calculate the upper bound $\hat{\eta}^+$ by solving (14), and the lower bound $\hat{\eta}^-$ by (15).

Output: $\hat{\eta}^+, \hat{\eta}^-$.

Proposition 4.1. Let $\hat{\eta}(Q)$ and $\hat{L}_K(Q)$ be the estimators of $\eta(Q)$ and $L_K(Q)$, such that for $\delta \in (0, 1)$,

$$\begin{aligned} \mathbb{P}[\eta(Q^\pi) \leq \hat{\eta}(Q^\pi) + \lambda_\eta] &\leq 1 - \delta, \\ \mathbb{P}[\hat{L}_K(Q^\pi) \leq \lambda_K] &\leq 1 - \delta, \end{aligned} \quad (9)$$

where λ_η and λ_K are two constants. Note that (9) only needs to hold for $Q = Q^\pi$. Assume $Q^\pi \in \mathcal{F}$. Define

$$\hat{\eta}^+ := \max_{Q \in \mathcal{F}} \left\{ [\hat{\eta}(Q) + \lambda_\eta], \text{ s.t. } \hat{L}_K(Q) \leq \lambda_K \right\}. \quad (10)$$

Then we have $\hat{\eta}^+ \geq \eta^\pi$ with probability $1 - 2\delta$.

Proof. From the assumption, with probability at least $1 - 2\delta$, we have both $\eta(Q^\pi) \leq \hat{\eta}(Q^\pi) + \lambda_\eta$ and $\hat{L}_K(Q^\pi) \leq \lambda_K$. In this case, Q^π belongs to the feasible set of the optimization in (10), i.e., $Q^\pi \in \mathcal{F}$, $\hat{L}_K(Q^\pi) \leq \lambda_K$. Therefore, we have $\eta^\pi = \eta(Q^\pi) \leq \hat{\eta}(Q^\pi) + \lambda_\eta \leq \hat{\eta}^+$ by the definition of the max operator. \square

We can use the kernel Bellman V-statistics in Proposition 3.1 to construct $\hat{L}_K(Q^\pi)$ and λ_K . For $\eta(Q^\pi)$, we use the Monte Carlo estimator $\hat{\eta}(Q^\pi)$ in (3) and set $\lambda_\eta = Q_{\max} \sqrt{2 \log(2/\delta)/N}$ by Hoeffding inequality.

4.2. Optimization in Reproducing Kernel Hilbert Space

To provide tight high confidence bounds, we need to choose \mathcal{F} to be a function space that is both simple and rich enough to contain Q^π . Here we choose \mathcal{F} to be a reproducing kernel Hilbert space (RKHS) \mathcal{H}_{K_0} induced by a positive kernel $K_0(x, \bar{x})$. We should distinguish K_0 with the kernel K used in kernel Bellman loss. Using RKHS allows us to incorporate a rich, infinite dimensional set of functions, and still obtain computationally tractable solution for the optimization in (10).

Proposition 4.2. For RKHS \mathcal{H}_{K_0} with kernel K_0 , define

$$\mathcal{F} = \{Q \in \mathcal{H}_{K_0} : \|Q\|_{\mathcal{H}_{K_0}}^2 \leq \rho\}, \quad (11)$$

where ρ is a positive constant. Using the \mathcal{F} in (11), the optimization solution of (10) can be expressed as

$$Q(\cdot) = \sum_{i=0}^n \alpha_i f_i(\cdot). \quad (12)$$

Here $\alpha = [\alpha_i]_{i=0}^n$ are the coefficients decided by the optimization problem, and $f_0(\cdot) = \mathbb{E}_{x \sim \mu_0 \times \pi} [K_0(\cdot, x)]$,

$$f_i(\cdot) = K_0(\cdot, x_i) - \gamma \mathbb{E}_{a'_i \sim \pi(\cdot | s'_i)} [K_0(\cdot, x'_i)], \quad i = 1, \dots, n,$$

where $x'_i = (s'_i, a'_i)$, with s'_i the observed next state following $x_i = (s_i, a_i)$ and a'_i randomly drawn from $\pi(\cdot | s'_i)$.

In addition, for Q of form (12), $\hat{\eta}(Q)$ is a linear function of α , and both $\hat{L}_K(Q)$ and $\|Q\|_{\mathcal{H}_{K_0}}^2$ are convex quadratic functions of α . In particular, we have $\|Q\|_{\mathcal{H}_{K_0}}^2 = \alpha^\top B \alpha$, where $B = [B_{ij}]$ is a $(n+1) \times (n+1)$ matrix with $B_{ij} = \langle f_i, f_j \rangle_{\mathcal{H}_{K_0}}$. Therefore, the optimization in (10) reduces to an optimization on α with linear objective and convex quadratic constraints,

$$\max_{\alpha^\top B \alpha \leq \rho^2} \left\{ [c^\top \alpha + \lambda_\eta], \text{ s.t. } (\alpha - b)^\top A (\alpha - b) \leq \lambda_K \right\},$$

where A, B are two $(n+1) \times (n+1)$ positive definite matrices and b, c two $(n+1) \times 1$ vectors, whose definition can be found in Appendix E.

Random Feature Approximation Unfortunately, solving the programming in Proposition 4.2 requires an $O((n+1)^3)$ time complexity and is too slow when the data size n is very large. We can speed up the computation by using random feature approximations. The idea is that any positive definite kernel can be expressed as $K_0(x, y) = \int_{\mathcal{W}} \phi(x, w) \phi(y, w) d\mu(w)$, where $\phi(x, w)$ denotes a feature map indexed by a parameter w in some space \mathcal{W} and μ is a measure on \mathcal{W} . A typical example of this is the random Fourier expansion of stationary kernels by Bochner's theorem (Rahimi & Recht, 2007), in which $\phi(x, w) = \cos(w^\top [x, 1])$.

To speed up the computation, we draw i.i.d. $\{w_i\}_{i=1}^n$ from μ and take

$$K_0(x, y) = \frac{1}{m} \sum_{k=1}^m \phi(x, w_k) \phi(y, w_k). \quad (13)$$

Then one can show that any function Q in the RKHS of K_0 in (13) can be expressed as

$$Q(x) = \theta^\top \Phi(x),$$

Algorithm 2 Post-hoc Diagnosis for Existing Estimators

Input: Off-policy data $\mathcal{D} = \{x_i, r_i, s'_i\}_{i \leq n}$; maximum reward r_{\max} , discounted factor γ , positive definite kernel K , RKHS norm ρ , random feature $\Phi(\cdot)$, fail probability δ , Monte Carlo sample size N , existing estimation \widehat{Q} .

Draw $\{x_0^i\}_{i=1}^N$ from the initial distribution $\mu_0 \times \pi$; Decide λ_η and λ_K by concentration inequalities.

Calculate the upper bound $\widehat{\eta}^+$ by solving (16), the lower bound $\widehat{\eta}^-$ by (17), and the debias function Q_{debias} by (18).

Output: The upper and lower bounds $\widehat{\eta}^+$, $\widehat{\eta}^-$, and $Q_{\text{debias}}(x) = \theta^\top \Phi(x)$.

where $\Phi(x) = [\phi(x, w_i)]_{i=1}^m$, with $\|Q\|_{\mathcal{H}_{K_0}}^2 = m \|\theta\|_2^2$, where $\|\cdot\|_2$ is the typical L_2 norm on \mathbb{R}^m .

Given the off-policy dataset $\mathcal{D} = \{x_i, r_i, s'_i\}$ with $x_i = (s_i, a_i)$, and an i.i.d. sample $\{x_{i,0}\}_{i=1}^N$ with $x_{i,0} = (s_{i,0}, a_{i,0})$ from the initial distribution $\mu_0 \times \pi$. The optimization in (10) can be shown to reduce to

$$\widehat{\eta}^+ = \max_{\|\theta\|_2^2 \leq \rho/m} \left\{ [c_0^\top \theta + \lambda_\eta], \right. \\ \left. \text{s.t. } (Z\theta - v)^\top M (Z\theta - v) \leq \lambda_K \right\}, \quad (14)$$

where $v = [r_i]_{i=1}^n \in \mathbb{R}^{n \times 1}$, $M \in \mathbb{R}^{n \times n}$ with $M_{ij} = [K(x_i, x_j)/n^2]_{ij}$, $c_0 = \sum_{i=1}^N \Phi(x_{i,0})/N \in \mathbb{R}^{m \times 1}$, and

$$Z = [\Phi(x_i) - \gamma \mathbb{E}_{a'_i \sim \pi(\cdot|s'_i)} [\Phi(x'_i)]]_{i=1}^n \in \mathbb{R}^{n \times m},$$

with $x'_i = [s'_i, a'_i]$. The expectation in Z can be approximated by Monte Carlo sampling from $\pi(\cdot|s'_i)$.

Similarly, we can get the lower confidence bounds via

$$\widehat{\eta}^- = \min_{\|\theta\|_2^2 \leq \rho/m} \left\{ [c_0^\top \theta - \lambda_\eta], \right. \\ \left. \text{s.t. } (Z\theta - v)^\top M (Z\theta - v) \leq \lambda_K \right\}. \quad (15)$$

Compared with the programming in Proposition (4.2), the optimization problems in (14)-(15) have lower dimensions and is hence much faster when $m \ll n$, since the dimension of θ is m , while that of α is $n + 1$. We describe the detailed procedure for obtaining the upper and lower confidence bounds in Algorithm 1.

4.3. Post-hoc Confidence Bound of Existing Estimators

We extend our method to provide post-hoc confidence bounds around existing estimators provided by users. Given an existing estimator \widehat{Q} of the Q-function, we denote by $Q_{\text{res}} := Q^\pi - \widehat{Q}$ the difference between the ground truth Q^π and the prior estimate \widehat{Q} . Assuming Q_{res} belongs to \mathcal{F} , we obtain an upper bound by

$$\max_{Q_{\text{res}} \in \mathcal{F}} \left\{ [\widehat{\eta}(Q_{\text{res}} + \widehat{Q}) + \lambda_\eta], \text{ s.t. } \widehat{L}_K(\widehat{Q} + Q_{\text{res}}) \leq \lambda_K \right\}.$$

This can be viewed as a special case of (10) but with the function space anchored around the existing estimator \widehat{Q} .

Similar to the earlier case, in the case of random feature approximation, the optimization reduces to

$$\widehat{\eta}^+ = \max_{\|\theta\|_2^2 \leq \rho/m} \left\{ [\widehat{\eta}(\widehat{Q}) + c_0^\top \theta + \lambda_\eta], \right. \\ \left. \text{s.t. } (Z\theta - \zeta)^\top M (Z\theta - \zeta) \leq \lambda_K \right\}, \quad (16)$$

where c_0, Z, M are defined as before, $\widehat{\eta}(\widehat{Q}) = \sum_{i=1}^N \widehat{Q}(x_{i,0})/N$, and $\zeta \in \mathbb{R}^{n \times 1}$ is defined to be the TD error vector of $\widehat{Q}(x)$ evaluated at the dataset \mathcal{D} ; that is, $\zeta = [\zeta_i]_{i=1}^n$ with

$$\zeta_i = r_i + \gamma \mathbb{E}_{a'_i \sim \pi(\cdot|s'_i)} \widehat{Q}(s'_i, a'_i) - \widehat{Q}(x_i).$$

The post-hoc lower bound follows a similar form:

$$\widehat{\eta}^- = \min_{\|\theta\|_2^2 \leq \rho/m} \left\{ [\widehat{\eta}(\widehat{Q}) + c_0^\top \theta - \lambda_\eta], \right. \\ \left. \text{s.t. } (Z\theta - \zeta)^\top M (Z\theta - \zeta) \leq \lambda_K \right\}. \quad (17)$$

4.4. Post-hoc Correction of Existing Estimators

In addition to providing confidence bounds around the existing estimator \widehat{Q} , we may want to further correct \widehat{Q} when \widehat{Q} is identified to admit large error. The post-hoc correction should ensure the corrected estimation falls into the confidence bounds provided earlier, while keeping the information in \widehat{Q} as much as possible.

Our idea is to correct \widehat{Q} by adding a debiasing term Q_{debias} , such that $\widehat{L}_K^V(\widehat{Q} + Q_{\text{debias}}) \leq \lambda_K$, while keeping Q_{debias} as small as possible. This is framed into

$$\min_{Q_{\text{debias}} \in \mathcal{F}} \left\{ \|Q_{\text{debias}}\|_{\mathcal{H}_{K_0}}^2 \text{ s.t. } \widehat{L}_K(Q + Q_{\text{debias}}) \leq \lambda_K \right\}.$$

We should distinguish Q_{debias} with the Q_{res} in Section 4.3, and Q_{res} can not be used for the debiasing purpose because it is designed to give an extreme estimation (for providing the bound), rather than a minimum correction.

In the case of random features approximation, the optimization reduces to

$$\theta^* = \arg \min_{\theta} \left\{ \|\theta\|_2^2, \right. \\ \left. \text{s.t. } (Z\theta - \zeta)^\top M (Z\theta - \zeta) \leq \lambda_K \right\}, \quad (18)$$




Figure 1. Results on Inverted-Pendulum (a)-(d) and Puck-Mountain (e)-(h). The plots show the confidence bounds on the average discounted reward as we vary the number of transitions n in (a) & (e), failure probability δ in (b) & (f), the temperature parameter τ of the behavior policy in (c) & (g), and the number of features in (d) & (h). The default parameters (when not varied) are: discounted factor $\gamma = 0.95$; horizon length $T = 50$ for Inverted-Pendulum and $T = 100$ for Puck-Mountain; number of episodes 20; failure probability $\delta = 0.10$; temperature of the behavior policy $\tau = 1$; and the feature dimension 10.

and $Q_{\text{debias}}(x) = \theta^{*\top} \Phi(x)$. If the existing estimator \hat{Q} already satisfies $\hat{L}_K^V(\hat{Q}) \leq \lambda_K$, this provides a zero correction (i.e., $\theta^* = 0$), since the estimator is already sufficiently accurate. This procedure is summarized in Algorithm 2.

5. Experiments

We evaluate the proposed algorithms in Section 4 on two continuous control tasks: Inverted-Pendulum and Puck-Mountain. Details of the tasks are in Appendix D.

For all of our experiments, we use a Gaussian RBF kernel $K(x_i, x_j) = \exp(-\|x_i - x_j\|_2^2 / h^2)$ for evaluating the kernel Bellman statistics, with a bandwidth selected from a separate batch of training data, and V -statistics to calculate Equation (5) given a set of empirical data.

To construct the behavior and target policies, we first train an optimal Q-function using deep Q-learning, and use its softmax functions as policies, and set the temperature to be higher for the behavior policies to encourage exploration. Note that our algorithms do not require to know the behavior policies since they are behavior-agnostic. A description for how to construct policies is in Appendix D.

5.1. Confidence Bounds for Off-Policy Evaluation

We test Algorithm 1 on the two continuous control tasks, Inverted-Pendulum (Figure 1(a)-(d)) and Puck-Mountain (Figure 1(e)-(h)). We solve the convex optimization in Algorithm 1 using CVXPY (Diamond & Boyd, 2016; Agrawal et al., 2018), which gives us the upper and lower bound $\hat{\eta}^+$ and $\hat{\eta}^-$. The results are reported w.r.t. the average discounted reward defined by $\eta^\pi / (1 - \gamma)$. The results are averaged over 50 random trials.

We use two types of feature maps $\Phi(\cdot)$ to parameterize the state-action value function $Q(x) := \theta^\top \Phi(x)$:

1) Random Fourier features: $\Phi(x) := [\cos(\mu_i^\top x + b_i)]_{i=1}^m$, where $\mu_i \sim \mathcal{N}(0, \frac{1}{h_0^2} \mathbf{I})$, $b_i \sim \text{Uniform}[0, 2\pi]$, and h_0 is a bandwidth parameter.

2) Neural features. We use a small neural network to parameterize Q function, and learn the Q function by minimizing the kernel loss on the training dataset, and set $\Phi(\cdot)$ by selecting a set of neural features (the neural feature map before the last linear layer) on the validation set.

Figure 1(a)-(d) demonstrate the evaluation results on Inverted-Pendulum. From Figure 1(a) & (b) we can see that, as we increase the number of transitions, or increase the failure probability δ , the confidence bounds become




Figure 2. Post-hoc diagnosis on Inverted-Pendulum. In figure (a)-(c), we vary the number of transitions with fixed $\delta = 0.10$. In figure (d)-(f), we fixed number of transitions $n = 500$ and vary the failure probability δ . The default parameters are: discounted factor $\gamma = 0.95$, the horizon length $T = 50$, number of transitions $n = 500$, failure probability $\delta = 0.10$, temperature of the behavior policy $\tau = 1$, and the feature dimension 10.

tighter since λ_K becomes smaller. However, it still covers the ground truth η^π , which indicates that our method gives a reasonable confidence interval for off-policy estimation.

In Figure 1(c), we investigate the performance of our method when we vary a temperature parameter of the behavior policy that controls the concentration of the probability mass. The confidence bounds do not change significantly with different temperature of behavior policies. In Figure 1(d), we study the performance of our algorithms with different number of features (including both random Fourier and neural features), which shows that we can get a tight confidence interval when the number of features is small. This is because when decreasing the number of features, the Q function is constrained in a lower dimensional function space and hence gives a tighter confidence interval. However, decreasing the number of features also increases the risk of model misspecification. We also test our method on Puck-Mountain, and report the results in Figure 1 (e)-(h), which show similar results as we observe in Inverted-Pendulum.

5.2. Post-hoc Diagnosis for Existing Estimators

We implement Algorithm 2 to provide post-hoc diagnosis for existing Q estimators, and test it on Inverted-Pendulum (Figure 2) and Puck-Mountain (Figure 3 in the Appendix D). Figure 2 (a)-(f) show the diagnosis results for two different state-action value function estimators (\hat{Q}_1 and \hat{Q}_2) on Inverted-Pendulum, which we learn with kernel Bellman statistics using different number of optimization steps. Here \hat{Q}_1 is a relatively accurate estimator while the estimation of \hat{Q}_2 has larger bias.

Figure 2 (a)-(c) show that as we increase the number of transitions, the norm of the debias term Q_{debias} keeps zero when \hat{Q}_1 and \hat{Q}_2 are inside the confidence interval. Meanwhile, when \hat{Q}_1 and \hat{Q}_2 can not provide an accurate estimation (falling outside of the confidence interval), our algorithm gives a good post-hoc correction to ensure the corrected estimation are inside the confidence interval. As we can see, such post-hoc diagnosis provides both confidence bounds and correction for existing estimators, which can be useful in real-world applications.

Figure 2 (d)-(f) demonstrate the performance of our algo-

rithms when we change the failure probability δ . Again, the empirical results shows that our algorithms can provide meaningful post-hoc diagnosis.

We also test our method on Puck-Mountain, following the same procedure as on Inverted-Pendulum, which shows similar results as we find in Inverted-Pendulum.

6. Conclusion

In this paper, we develop a new variational framework for constructing confidence bounds for off-policy evaluation and extend it to perform post-hoc diagnosis on existing estimators. In future work, we will leverage our framework to develop safe and efficient policy optimization and exploration algorithms based on the kernel Bellman statistics.

Acknowledgment

This work is supported in part by NSF CRII 1830161 and NSF CAREER 1846421.

References

- Agrawal, A., Verschueren, R., Diamond, S., and Boyd, S. A rewriting system for convex optimization problems. *Journal of Control and Decision*, 5(1):42–60, 2018.
- Baird, L. C. Residual algorithms: Reinforcement learning with function approximation. In *Proceedings of the 12th International Conference on Machine Learning*, pp. 30–37, 1995.
- Bottou, L., Peters, J., Quiñero-Candela, J., Charles, D. X., Chickering, D. M., Portugaly, E., Ray, D., Simard, P., and Snelson, E. Counterfactual reasoning and learning systems: The example of computational advertising. *Journal of Machine Learning Research*, 14:3207–3260, 2013.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- Dann, C., Li, L., Wei, W., and Brunskill, E. Policy certificates: Towards accountable reinforcement learning. In *Proceedings of the 36th International Conference on Machine Learning*, pp. 1507–1516, 2019.
- Diamond, S. and Boyd, S. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.
- Duan, Y., Jia, Z., and Wang, M. Minimax-optimal off-policy evaluation with linear function approximation. *Proceedings of the 37th International Conference on Machine Learning*, 2020.
- Dudík, M., Langford, J., and Li, L. Doubly robust policy evaluation and learning. In *Proceedings of the 28th International Conference on Machine Learning*, pp. 1097–1104, 2011.
- Farahmand, A.-m., Ghavamzadeh, M., Szepesvári, C., and Mannor, S. Regularized policy iteration with nonparametric function spaces. *The Journal of Machine Learning Research*, 17(1):4809–4874, 2016.
- Farajtabar, M., Chow, Y., and Ghavamzadeh, M. More robust doubly robust off-policy evaluation. In *Proceedings of the 35th International Conference on Machine Learning*, pp. 1446–1455, 2018.
- Feng, Y., Li, L., and Liu, Q. A kernel loss for solving the bellman equation. In *Advances in Neural Information Processing Systems*, pp. 15456–15467, 2019.
- Grunewalder, S., Lever, G., Baldassarre, L., Pontil, M., and Gretton, A. Modelling transition dynamics in mdps with rkhs embeddings. *Proceedings of the 29th International Conference on Machine Learning*, 2012.
- Han, F. An exponential inequality for u-statistics under mixing conditions. *Journal of Theoretical Probability*, 31(1):556–578, 2018.
- Hanna, J. P., Stone, P., and Niekum, S. Bootstrapping with models: Confidence intervals for off-policy evaluation. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- Hirano, K., Imbens, G. W., and Ridder, G. Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica*, 71(4):1161–1189, 2003.
- Hoeffding, W. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
- Jiang, N. and Huang, J. Minimax confidence interval for off-policy evaluation and policy optimization. *arXiv preprint arXiv:2002.02081*, 2020.
- Jiang, N. and Li, L. Doubly robust off-policy evaluation for reinforcement learning. In *Proceedings of the 23rd International Conference on Machine Learning*, pp. 652–661, 2016.
- Jin, C., Allen-Zhu, Z., Bubeck, S., and Jordan, M. I. Is q-learning provably efficient? In *Advances in Neural Information Processing Systems*, pp. 4863–4873, 2018.
- Jin, C., Yang, Z., Wang, Z., and Jordan, M. I. Provably efficient reinforcement learning with linear function approximation. *Proceedings of the 33rd Conference On Learning Theory*, 2020.

- Jung, T. and Polani, D. Kernelizing Ispe (λ). In *2007 IEEE International Symposium on Approximate Dynamic Programming and Reinforcement Learning*, pp. 338–345. IEEE, 2007.
- Kallus, N. and Uehara, M. Double reinforcement learning for efficient and robust off-policy evaluation. *Proceedings of the 37th International Conference on Machine Learning*, 2020.
- Li, L., Chu, W., Langford, J., and Wang, X. Unbiased offline evaluation of contextual-bandit-based news article recommendation algorithms. In *Proceedings of the 4th International Conference on Web Search and Data Mining (WSDM)*, pp. 297–306, 2011.
- Liao, P., Klasnja, P., and Murphy, S. Off-policy estimation of long-term average outcomes with applications to mobile health. *arXiv preprint arXiv:1912.13088*, 2019.
- Liu, Q., Li, L., Tang, Z., and Zhou, D. Breaking the curse of horizon: Infinite-horizon off-policy estimation. In *Advances in Neural Information Processing Systems*, pp. 5356–5366, 2018.
- Mousavi, A., Li, L., Liu, Q., and Zhou, D. Black-box off-policy estimation for infinite-horizon reinforcement learning. In *International Conference on Learning Representations*, 2020.
- Murphy, S. A., van der Laan, M., and Robins, J. M. Marginal mean models for dynamic regimes. *Journal of the American Statistical Association*, 96(456):1410–1423, 2001.
- Nachum, O., Chow, Y., Dai, B., and Li, L. Dualdice: Behavior-agnostic estimation of discounted stationary distribution corrections. In *Advances in Neural Information Processing Systems*, pp. 2318–2328, 2019.
- Puterman, M. L. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley-Interscience, New York, 1994. ISBN 0-471-61977-9.
- Rahimi, A. and Recht, B. Random features for large-scale kernel machines. In *Advances in neural information processing systems*, pp. 1177–1184, 2007.
- Sutton, R. S. and Barto, A. G. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, March 1998. ISBN 0-262-19398-1.
- Tang, Z., Feng, Y., Li, L., Zhou, D., and Liu, Q. Doubly robust bias reduction in infinite horizon off-policy estimation. In *International Conference on Learning Representations*, 2020.
- Taylor, G. and Parr, R. Kernelized value function approximation for reinforcement learning. In *Proceedings of the Twenty-Sixth International Conference on Machine Learning*, pp. 1017–1024, 2009.
- Thomas, P. S. and Brunskill, E. Data-efficient off-policy policy evaluation for reinforcement learning. In *Proceedings of the 33rd International Conference on Machine Learning*, pp. 2139–2148, 2016.
- Thomas, P. S., Theocharous, G., and Ghavamzadeh, M. High confidence policy improvement. In *Proceedings of the 32nd International Conference on Machine Learning*, pp. 2380–2388, 2015a.
- Thomas, P. S., Theocharous, G., and Ghavamzadeh, M. High-confidence off-policy evaluation. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015b.
- Thomas, P. S., Theocharous, G., Ghavamzadeh, M., Durgar, I., and Brunskill, E. Predictive off-policy policy evaluation for nonstationary decision problems, with applications to digital marketing. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, pp. 4740–4745, 2017.
- Uehara, M., Huang, J., and Jiang, N. Minimax weight and q-function learning for off-policy evaluation. *Proceedings of the 37th International Conference on Machine Learning*, 2020.
- Wang, Y.-X., Agarwal, A., and Dudík, M. Optimal and adaptive off-policy evaluation in contextual bandits. In *Proceedings of the 34th International Conference on Machine Learning*, pp. 3589–3597, 2017.
- White, M. and White, A. Interval estimation for reinforcement-learning algorithms in continuous-state domains. In *Advances in Neural Information Processing Systems*, pp. 2433–2441, 2010.
- Xie, T., Ma, Y., and Wang, Y.-X. Towards optimal off-policy evaluation for reinforcement learning with marginalized importance sampling. In *Advances in Neural Information Processing Systems*, pp. 9668–9678, 2019.
- Xu, X., Xie, T., Hu, D., and Lu, X. Kernel least-squares temporal difference learning. *International Journal of Information and Technology*, 11(9):54–63, 2005.
- Xu, X., Hu, D., and Lu, X. Kernel-based least-squares policy iteration for reinforcement learning. *IEEE Transactions on Neural Networks*, 18(4):973–992, 2007.
- Zhang, R., Dai, B., Li, L., and Schuurmans, D. Gendice: Generalized offline estimation of stationary values. In *International Conference on Learning Representations*, 2020.