

Efficient Data Reuse in Value Function Approximation.

Hirotaka Hachiya[†] and Takayuki Akiyama[†] and Masashi Sugiyama[†] and Jan Peters[‡]

Abstract—Off-policy reinforcement learning is aimed at efficiently using data samples gathered from a policy that is different from the currently optimized policy. A common approach is to use importance sampling techniques for compensating for the bias of value function estimators caused by the difference between the data-sampling policy and the target policy. However, existing off-policy methods often do not take the variance of the value function estimators explicitly into account and therefore their performance tends to be unstable. To cope with this problem, we propose using an adaptive importance sampling technique which allows us to actively control the trade-off between bias and variance. We further provide a method for optimally determining the trade-off parameter based on a variant of cross-validation. The usefulness of the proposed approach is demonstrated through simulated swing-up inverted-pendulum problem.

I. INTRODUCTION

Policy iteration is a reinforcement learning setup where the optimal policy is obtained by iteratively performing policy evaluation and improvement steps [17], [2]. When policies are updated, many popular policy iteration methods require the user to gather new samples following the updated policy, and the new samples are used for *value function approximation*. However, this approach is inefficient particularly when the sampling cost is high and it would be more cost-efficient if we could reuse the data collected in the past. A situation where the sampling policy (a policy used for gathering data samples) and the current policy are different is called *off-policy* reinforcement learning [17].

In the off-policy setup, simply employing a standard policy iteration method such as *least-squares* policy iteration [6] does not lead to the optimal policy as the sampling policy can introduce bias into value function approximation. This distribution mismatch problem can be eased by the use of *importance sampling* techniques [4], which cancel the bias asymptotically. However, the approximation error is not necessarily small when the bias is reduced by importance sampling; the variance of estimators also needs to be taken into account since the approximation error is the sum of squared bias and variance. Due to large variance, existing importance sampling techniques tend to be unstable [17], [11].

[†]Department of Computer Science, Tokyo Institute of Technology, 2-12-1 O-okayama, Meguro-ku, Tokyo 152-8552, Japan (email: {hachiya@sg, akiyama@sg, sugi@}cs.titech.ac.jp), [‡]Max-Planck Institute for Biological Cybernetics, Dept. Schölkopf, Spemannstraße 38, Rm 223, 72076 Tübingen, Germany (email: jan.peters@tuebingen.mpg.de)

This research was supported in part by JSPS Global COE program “Computationism as a Foundation for the Sciences” and MEXT (20680007 and 18300057).

To overcome the instability problem, we propose using an *adaptive importance sampling* technique used in statistics [15]. The proposed adaptive method, which smoothly bridges the ordinary estimator and importance-weighted estimator, allows us to control the trade-off between bias and variance. Thus, given that the trade-off parameter is determined carefully, the optimal performance can be achieved in terms of both bias and variance. However, the optimal value of the trade-off parameter is heavily dependent on data samples and policies, and therefore using a pre-determined parameter value may not be always effective in practice.

For optimally choosing the value of the trade-off parameter, we propose using an automatic model selection method based on a variant of cross-validation [16]. The method called *importance-weighted cross-validation* enables us to estimate the approximation error of value functions in an almost unbiased manner even under off-policy situations. Thus we can adaptively choose the trade-off parameter based on data samples at hand. The usefulness of the proposed approach is demonstrated through simulated swing-up inverted-pendulum problem.

II. BACKGROUND AND NOTATION

In this section, we review how Markov decision problems can be solved using policy iteration based on value functions.

A. Markov Decision Problems

Let us consider a Markov decision problem (MDP) specified by $(\mathcal{S}, \mathcal{A}, P_{\mathbb{T}}, R, \gamma)$, where \mathcal{S} is a set of states, \mathcal{A} is a set of actions, $P_{\mathbb{T}}(s'|s, a) \in [0, 1]$ is the transition probability-density from state s to next state s' when action a is taken, $R(s, a, s') \in \mathbb{R}$ is a reward for transition from s to s' by taking action a , $\gamma \in (0, 1]$ is the discount factor for future rewards. Let $\pi(a|s) \in [0, 1]$ be a stochastic policy which is the conditional probability density of taking action a given state s . The state-action value function $Q^{\pi}(s, a) \in \mathbb{R}$ for policy π is the expected discounted sum of rewards the agent will receive when taking action a in state s and following policy π thereafter, i.e.,

$$Q^{\pi}(s, a) \equiv \mathbb{E}_{\pi, P_{\mathbb{T}}} \left[\sum_{n=1}^{\infty} \gamma^{n-1} R(s_n, a_n, s_{n+1}) \mid s_1 = s, a_1 = a \right],$$

where $\mathbb{E}_{\pi, P_{\mathbb{T}}}$ denotes the expectation over $\{s_n, a_n\}_{n=1}^{\infty}$ following $\pi(a_n|s_n)$ and $P_{\mathbb{T}}(s_{n+1}|s_n, a_n)$.

The goal of reinforcement learning is to obtain the policy which maximizes the sum of future rewards; the optimal

policy can be expressed¹ as

$$\pi^*(a|s) \equiv \delta(a - \arg \max_{a'} Q^*(s, a')),$$

where $\delta(\cdot)$ is the Dirac delta function and $Q^*(s, a)$ is the *optimal* state-action value function defined by

$$Q^*(s, a) \equiv \max_{\pi} Q^{\pi}(s, a).$$

$Q^{\pi}(s, a)$ can be expressed as the following recurrent form called the *Bellman equation* [17]:

$$Q^{\pi}(s, a) = R(s, a) + \gamma \mathbb{E}_{P_{\mathcal{T}}(s'|s, a)} \mathbb{E}_{\pi(a'|s')} [Q^{\pi}(s', a')], \quad (1)$$

where $R(s, a)$ is the expected reward function when the agent takes action a in state s :

$$R(s, a) \equiv \mathbb{E}_{P_{\mathcal{T}}(s'|s, a)} [R(s, a, s')].$$

$\mathbb{E}_{P_{\mathcal{T}}(s'|s, a)}$ denotes the conditional expectation of s' over $P_{\mathcal{T}}(s'|s, a)$ given s and a . $\mathbb{E}_{\pi(a'|s')}$ denotes the conditional expectation of a' over $\pi(a'|s')$ given s' .

B. Policy Iteration

The computation of the value function $Q^{\pi}(s, a)$ is called *policy evaluation*. Using $Q^{\pi}(s, a)$, we can find a better policy $\pi'(a|s)$ by

$$\pi'(a|s) = \delta(a - \arg \max_{a'} Q^{\pi}(s, a')).$$

This is called (greedy) *policy improvement*. It is known that repeating policy evaluation and policy improvement results in the optimal policy $\pi^*(a|s)$ [17]. This entire process is called *policy iteration*:

$$\pi_1 \xrightarrow{E} Q^{\pi_1} \xrightarrow{I} \pi_2 \xrightarrow{E} Q^{\pi_2} \xrightarrow{I} \pi_3 \xrightarrow{E} \dots \xrightarrow{I} \pi^*,$$

where π_1 is an initial policy, E and I indicate policy *evaluation* and *improvement* steps respectively. For technical reasons, we assume that all policies are strictly positive (i.e., all actions have non-zero probability densities). In order to guarantee this, we use policy improvement which generates explorative policies such as the *Gibbs policy* and the ϵ -*greedy policy*. In the case of the Gibbs policy,

$$\pi'(a|s) = \frac{\exp(Q^{\pi}(s, a)/\tau)}{\int_{\mathcal{A}} \exp(Q^{\pi}(s, a')/\tau) da'}, \quad (2)$$

where τ is a positive parameter which determines the randomness of the new policy π' . In the case of the ϵ -greedy policy,

$$\pi'(a|s) = \begin{cases} 1 - \epsilon + \epsilon/|\mathcal{A}| & \text{if } a = a^*, \\ \epsilon/|\mathcal{A}| & \text{otherwise,} \end{cases} \quad (3)$$

where

$$a^* = \arg \max_a Q^{\pi}(s, a),$$

and $\epsilon \in (0, 1]$ determines how stochastic the new policy π' is.

¹We assume that given state s there is only one action maximizing the optimal value function $Q^*(s, a)$.

C. Value Function Approximation

Although policy iteration is guaranteed to produce the optimal policy, it is often computationally intractable since the number of state-action pairs $|\mathcal{S}| \times |\mathcal{A}|$ is very large; $|\mathcal{S}|$ or $|\mathcal{A}|$ becomes infinite when the state space or action space is continuous. To overcome this problem, the authors of the references [17], [10], [6] proposed to approximate the state-action value function $Q^{\pi}(s, a)$ using the following linear model:

$$\widehat{Q}^{\pi}(s, a; \boldsymbol{\theta}) \equiv \sum_{b=1}^B \theta_b \phi_b(s, a) = \boldsymbol{\theta}^{\top} \boldsymbol{\phi}(s, a),$$

where

$$\boldsymbol{\phi}(s, a) = (\phi_1(s, a), \phi_2(s, a), \dots, \phi_B(s, a))^{\top}$$

are the fixed basis functions, \top denotes the transpose, B is the number of basis functions, and

$$\boldsymbol{\theta} = (\theta_1, \theta_2, \dots, \theta_B)^{\top}$$

are model parameters. Note that B is usually chosen to be much smaller than $|\mathcal{S}| \times |\mathcal{A}|$. For N -step transitions, we ideally want to learn the parameters $\boldsymbol{\theta}$ so that the approximation error is minimized:

$$\min_{\boldsymbol{\theta}} \mathbb{E}_{P_1, \pi, P_{\mathcal{T}}} \left[\frac{1}{N} \sum_{n=1}^N \left(\widehat{Q}^{\pi}(s_n, a_n; \boldsymbol{\theta}) - Q^{\pi}(s_n, a_n) \right)^2 \right],$$

where $\mathbb{E}_{P_1, \pi, P_{\mathcal{T}}}$ denotes the expectation over $\{s_n, a_n\}_{n=1}^N$ following the initial-state probability density $P_1(s_1)$, the policy $\pi(a_n|s_n)$, and the transition probability density $P_{\mathcal{T}}(s_{n+1}|s_n, a_n)$.

A fundamental problem of the above formulation is that the target function $Q^{\pi}(s, a)$ cannot be observed directly. To cope with this problem, we use the square of the Bellman residual [13], [6] as

$$\boldsymbol{\theta}^* \equiv \arg \min_{\boldsymbol{\theta}} G, \quad (4)$$

$$G \equiv \mathbb{E}_{P_1, \pi, P_{\mathcal{T}}} \left[\frac{1}{N} \sum_{n=1}^N g(s_n, a_n; \boldsymbol{\theta}) \right],$$

$$g(s, a; \boldsymbol{\theta}) \equiv \left(\widehat{Q}^{\pi}(s, a; \boldsymbol{\theta}) - R(s, a) - \gamma \mathbb{E}_{P_{\mathcal{T}}(s'|s, a)} \mathbb{E}_{\pi(a'|s')} \left[\widehat{Q}^{\pi}(s', a'; \boldsymbol{\theta}) \right] \right)^2,$$

where $g(s, a; \boldsymbol{\theta})$ is the approximation error for one step (s, a) derived from the Bellman equation (1).

D. On-policy vs. Off-policy

We suppose that a dataset consisting of M episodes of N steps is available. The agent initially starts from a randomly selected state s_1 following the initial-state probability density $P_1(s)$ and chooses an action based on a *sampling policy* $\tilde{\pi}(a_n|s_n)$. Then the agent makes a transition following $P_{\mathcal{T}}(s_{n+1}|s_n, a_n)$ and receives a reward r_n (=

$R(s_n, a_n, s_{n+1})$). This is repeated for N steps—thus the training data $\mathcal{D}^{\tilde{\pi}}$ is expressed as

$$\mathcal{D}^{\tilde{\pi}} \equiv \{d_m^{\tilde{\pi}}\}_{m=1}^M,$$

where each episodic sample $d_m^{\tilde{\pi}}$ consists of a set of 4-tuple elements as

$$d_m^{\tilde{\pi}} \equiv \{(s_{m,n}^{\tilde{\pi}}, a_{m,n}^{\tilde{\pi}}, r_{m,n}^{\tilde{\pi}}, s_{m,n+1}^{\tilde{\pi}})\}_{n=1}^N.$$

We use two types of policies which have different purposes: the *sampling policy* $\tilde{\pi}(a|s)$ for collecting data samples and the *current policy* $\pi(a|s)$ for computing the value function \hat{Q}^π . When $\tilde{\pi}(a|s)$ is equal to $\pi(a|s)$ (the situation called *on-policy*), just replacing the expectation contained in the error G by sample averages gives a *consistent* estimator (i.e., the estimated parameter converges to the optimal value as the number M of episodes goes to infinity):

$$\begin{aligned} \hat{\theta}_{\text{NIW}} &\equiv \arg \min_{\theta} \hat{G}_{\text{NIW}}, \\ \hat{G}_{\text{NIW}} &\equiv \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N \hat{g}_{m,n}, \\ \hat{g}_{m,n} &\equiv \hat{g}(s_{m,n}^{\tilde{\pi}}, a_{m,n}^{\tilde{\pi}}; \theta, \mathcal{D}^\pi), \end{aligned}$$

$$\begin{aligned} \hat{g}(s, a; \theta, \mathcal{D}) &\equiv \left(\hat{Q}^\pi(s, a; \theta) - \frac{1}{|\mathcal{D}(s,a)|} \sum_{r \in \mathcal{D}(s,a)} r \right. \\ &\quad \left. - \frac{\gamma}{|\mathcal{D}(s,a)|} \sum_{s' \in \mathcal{D}(s,a)} \pi(a'|s') \mathbb{E} [\hat{Q}^\pi(s', a'; \theta)] \right)^2, \end{aligned}$$

where $\mathcal{D}(s,a)$ is a set of 4-tuple elements containing state s and action a in the training data \mathcal{D} , and $\sum_{r \in \mathcal{D}(s,a)}$ and $\sum_{s' \in \mathcal{D}(s,a)}$ denote the summation over r and s' in the set $\mathcal{D}(s,a)$, respectively. Note that ‘NIW’ stands for ‘No Importance Weight’ (explained later).

However, $\tilde{\pi}(a|s)$ is usually different from $\pi(a|s)$ in reality since the current policy is updated in policy iteration. The situation where $\tilde{\pi}(a|s)$ is different from $\pi(a|s)$ is called *off-policy*. In the off-policy setup, $\hat{\theta}_{\text{NIW}}$ is no longer consistent. This inconsistency problem could be avoided by gathering new samples, i.e., when the current policy is updated, new samples are gathered following the updated policy and the new samples are used for policy evaluation. However, when the data sampling cost is high, this is not cost-efficient—it would be more efficient if we could reuse the previously gathered samples.

In the following sections, we address the issue of sample reuse in the off-policy setup.

III. IMPORTANCE-WEIGHTING TECHNIQUES

In this section, we review existing off-policy reinforcement learning techniques.

A. Importance Sampling

Importance sampling is a general technique for dealing with the off-policy situation. Suppose we have i.i.d. (*independent and identically distributed*) samples $\{x_m\}_{m=1}^M$ from a strictly positive probability density function $\tilde{P}(x)$. Using these samples, we would like to compute the expectation of a function $g(x)$ over another probability density function $P(x)$. A consistent approximation of the expectation is given by the *importance-weighted* average as follows:

$$\begin{aligned} \frac{1}{M} \sum_{m=1}^M g(x_m) \frac{P(x_m)}{\tilde{P}(x_m)} &\xrightarrow{M \rightarrow \infty} \mathbb{E}_{\tilde{P}(x)} \left[g(x) \frac{P(x)}{\tilde{P}(x)} \right] \\ &= \int g(x) \frac{P(x)}{\tilde{P}(x)} \tilde{P}(x) dx = \int g(x) P(x) dx = \mathbb{E}_{P(x)} [g(x)]. \end{aligned}$$

However, applying the importance sampling technique in off-policy reinforcement learning is not that straightforward since our training samples of state s and action a are not i.i.d. due to the sequential nature of MDPs. Below, we review existing importance-weighting techniques in MDPs.

B. Episodic Importance-weights

The reference [17] proposed the *episodic importance-weight* (EIW) method which utilizes the independence between episodes:

$$\begin{aligned} P(d, d') &= P(d)P(d') \\ &= P(s_1, a_1, \dots, s_N, a_N, s_{N+1}) P(s'_1, a'_1, \dots, s'_N, a'_N, s'_{N+1}). \end{aligned}$$

Based on the independence between episodes, the error G defined by Eq.(4) can be rewritten as

$$G = \mathbb{E}_{P_I, \tilde{\pi}, P_T} \left[\frac{1}{N} \sum_{n=1}^N g(s_n, a_n; \theta) w_N \right],$$

where

$$w_N \equiv \frac{P_\pi(d)}{P_{\tilde{\pi}}(d)}.$$

$P_\pi(d)$ and $P_{\tilde{\pi}}(d)$ are the probability densities of observing episodic data d under policy π and $\tilde{\pi}$:

$$\begin{aligned} P_\pi(d) &\equiv P_I(s_1) \prod_{n=1}^N \pi(a_n | s_n) P_T(s_{n+1} | s_n, a_n), \\ P_{\tilde{\pi}}(d) &\equiv P_I(s_1) \prod_{n=1}^N \tilde{\pi}(a_n | s_n) P_T(s_{n+1} | s_n, a_n). \end{aligned}$$

We note that the importance weights can be computed without explicitly knowing P_I and P_T since they are canceled out:

$$w_N = \frac{\prod_{n=1}^N \pi(a_n | s_n)}{\prod_{n=1}^N \tilde{\pi}(a_n | s_n)}.$$

Using the training data $\mathcal{D}^{\tilde{\pi}}$, we can construct a consistent estimator of G as

$$\hat{G}_{\text{EIW}} \equiv \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N \hat{g}_{m,n} \hat{w}_{m,n}, \quad (5)$$

where

$$\hat{w}_{m,N} \equiv \frac{\prod_{n'=1}^N \pi(a_{m,n'}^{\tilde{\pi}} | s_{m,n'}^{\tilde{\pi}})}{\prod_{n'=1}^N \tilde{\pi}(a_{m,n'}^{\tilde{\pi}} | s_{m,n'}^{\tilde{\pi}})}.$$

Based on this, the parameter θ is estimated by

$$\hat{\theta}_{\text{EIW}} \equiv \arg \min_{\theta} \hat{G}_{\text{EIW}}.$$

C. Per-decision Importance-weights

In the reference [11], a more efficient importance sampling technique called the *per-decision importance-weight* (PDIW) method was proposed. A crucial observation in PDIW is that the error at the n -th step does not depend on the samples after the n -th step, i.e., the error G can be rewritten as

$$G = \mathbb{E}_{P_1, \tilde{\pi}, P_T} \left[\frac{1}{N} \sum_{n=1}^N g(s_n, a_n; \theta) w_n \right].$$

Using the training data $\mathcal{D}^{\tilde{\pi}}$, we can construct a consistent estimator as follows (cf. Eq.(5))

$$\hat{G}_{\text{PDIW}} \equiv \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N \hat{g}_{m,n} \hat{w}_{m,n}. \quad (6)$$

$\hat{w}_{m,n}$ in Eq.(6) *only* contains the relevant terms up to the n -th step, while $\hat{w}_{m,N}$ in Eq.(5) includes all the terms until the end of the episode.

Based on this, the parameter θ is estimated by

$$\hat{\theta}_{\text{PDIW}} \equiv \arg \min_{\theta} \hat{G}_{\text{PDIW}}.$$

IV. ADAPTIVE PER-DECISION IMPORTANCE-WEIGHTS

The importance-weighted estimator $\hat{\theta}_{\text{PDIW}}$ (also $\hat{\theta}_{\text{EIW}}$) is guaranteed to be consistent. However, both are not *efficient* in the statistical sense [15], i.e., they do not have the smallest admissible variance². For this reason, $\hat{\theta}_{\text{PDIW}}$ can have large variance in finite sample cases and therefore learning with PDIW could be unstable in practice. In this section, we propose a new importance-weighting method that is more stable than the existing methods.

A. Definition

In order to improve the estimation accuracy, it is important to control the trade-off between consistency and efficiency (or similarly bias and variance) based on the training data. Here, we introduce a *flattening parameter* $\nu \in [0, 1]$ to control the trade-off by slightly ‘flattening’ the importance weights [15], [16]:

$$\hat{G}_{\text{APDIW}} \equiv \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N \hat{g}_{m,n} (\hat{w}_{m,n})^{\nu}, \quad (7)$$

where APDIW means Adaptive PDIW. Based on this, the parameter θ is estimated as follows:

$$\hat{\theta}_{\text{APDIW}} \equiv \arg \min_{\theta} \hat{G}_{\text{APDIW}}.$$

²More precisely, an estimator $\hat{\theta}$ is said to be efficient if it is unbiased and achieves the Cramér-Rao lower-bound [12].

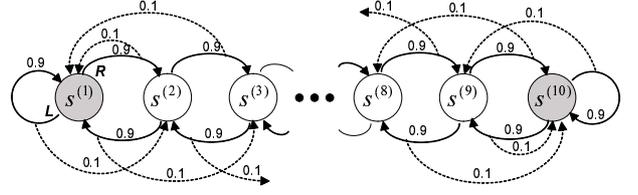


Fig. 1. 10-state chain walk MDP.

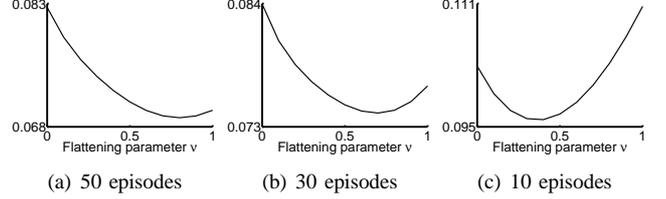


Fig. 2. Average true error G over 50 trials, as a function of the flattening parameter ν in the 10-state chain walk problem. The trend of G differs dependent on the number of episodes.

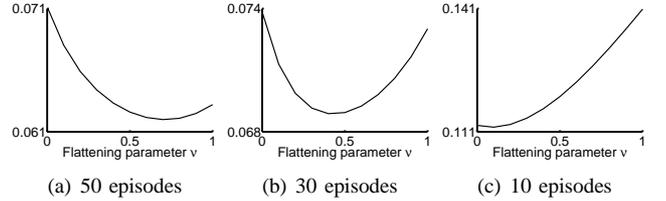


Fig. 3. Average error estimated by 5-fold IWCV (\hat{G}_{APDIW}), as a function of the flattening parameter ν in the 10-state chain walk problem. IWCV nicely captures the trend of the true error G in the Fig.2

When $\nu = 0$, $\hat{\theta}_{\text{APDIW}}$ is reduced to the ordinary estimator $\hat{\theta}_{\text{NIW}}$. Therefore, it has large bias but has relatively small variance. On the other hand, when $\nu = 1$, $\hat{\theta}_{\text{APDIW}}$ is reduced to the importance-weighted estimator $\hat{\theta}_{\text{PDIW}}$. Therefore, it has small bias but has relatively large variance. In practice, an intermediate ν will yield the best performance.

B. Numerical Examples

In order to illustrate how the flattening parameter ν influences the estimator $\hat{\theta}_{\text{APDIW}}$, we perform policy evaluation in the *chain-walk MDP* illustrated in Fig.1. The MDP consists of 10 states

$$\mathcal{S} = \{s^{(i)}\}_{i=1}^{10}$$

and two actions

$$\mathcal{A} = \{a^{(i)}\}_{i=1}^2 = \{L, R\}.$$

The reward +1 is given when visiting $s^{(1)}$ and $s^{(10)}$. The transition probability P_T is indicated by the numbers attached to the arrows in the figure; for example, $P_T(s^{(2)} | s^{(1)}, R) = 0.9$ and $P_T(s^{(1)} | s^{(1)}, R) = 0.1$, meaning that the agent can successfully move to the right node with probability 0.9 (indicated by solid arrows in the figure) and the action fails with probability 0.1 (indicated by dashed arrows in the figure). We use 6 Gaussian kernels with standard deviation $\sigma = 10$ as basis functions and locate kernel centers at $s^{(1)}$, $s^{(5)}$ and $s^{(10)}$. More specifically, the basis functions $\phi(s, a) = (\phi_1(s, a), \dots, \phi_6(s, a))$ are defined by

$$\phi_{3(i-1)+j}(s, a) = I(a = a^{(i)}) \exp\left(-\frac{(s - c_j)^2}{2\sigma^2}\right), \quad (8)$$

for $i = 1, 2$ and $j = 1, 2, 3$, where

$$c_1 = 1, c_2 = 5, c_3 = 10,$$

and

$$I(x) = \begin{cases} 1 & \text{if } x \text{ is true,} \\ 0 & \text{if } x \text{ is not true,} \end{cases}$$

We ran the experiments 50 times; the sampling policy $\tilde{\pi}(a|s)$ and the current policy $\pi(a|s)$ were chosen randomly at every trial such that $\tilde{\pi} \neq \pi$. The discount factor was set at $\gamma = 0.9$. The model parameter $\hat{\theta}_{\text{APDIW}}$ was learned from the training samples $\mathcal{D}_{\tilde{\pi}}$ and its approximation error was computed from the test samples \mathcal{D}_{π} .

Fig.2 depicts the true approximation error G averaged over 50 trials as a function of the flattening parameter ν for $M = 10, 30, 50$. Fig.2(a) shows that when the number of episodes is large ($M = 50$), the approximation error tends to decrease as the flattening parameter increases. This would be a natural result due to the consistency of $\hat{\theta}_{\text{APDIW}}$ when $\nu = 1$. On the other hand, Fig.2(b) shows that when the number of episodes is not large ($M = 30$), $\nu = 1$ performs rather poorly. This implies that the consistent estimator tends to be unstable when the number of episodes is not large enough— $\nu = 0.7$ works the best in this case. Fig.2(c) shows the results when the number of episodes is further reduced ($M = 10$). This illustrates that the consistent estimator with $\nu = 1$ is even worse than the ordinary estimator ($\nu = 0$) because the bias is dominated by large variance. In this case, the best ν is even smaller and is achieved at $\nu = 0.4$. The above results show that APDIW can outperform PDIW particularly when only a small number of training samples are available, provided that the flattening parameter ν is chosen appropriately.

V. AUTOMATIC SELECTION OF THE FLATTENING PARAMETER

Generally, the best ν tends to be large (small) when the number of training samples is large (small). However, this general trend is not sufficient to fine-tune the flattening parameter since the best value of ν depends on training samples, policies, the model of value functions etc. In this section, we discuss how we perform *model selection* in order to choose the best flattening parameter ν automatically from the training data and policies.

A. Importance-weighted Cross-validation

As shown in Section IV, the performance of APDIW depends on the choice of the flattening parameter ν . Ideally, we set ν so that the approximation error G is minimized, but true G is inaccessible in practice. To cope with this problem, we estimate the approximation error G using *importance-weighted cross-validation* (IWCV) [16]. The basic idea of IWCV is to divide the training data $\mathcal{D}_{\tilde{\pi}}$ into a ‘training part’ and a ‘validation part’. Then the parameter θ is learned from

the training part and its approximation error is estimated using the validation part. Below we explain in more detail how we apply IWCV to the selection of the flattening parameter ν in the current context.

Let us divide a training dataset $\mathcal{D}_{\tilde{\pi}}$ containing M episodes into K subsets $\{\mathcal{D}_{\tilde{\pi}_k}^k\}_{k=1}^K$ of approximately the same size (typically $K = 5$). For simplicity, we assume that M is divisible by K . Let $\hat{\theta}_{\text{APDIW}}^k$ be the parameter learned from $\{\mathcal{D}_{\tilde{\pi}_{k'}}^k\}_{k' \neq k}$ with APDIW (cf. Eq.(7)). Then, the approximation error is estimated by

$$\hat{G}_{\text{IWCV}} = \frac{1}{K} \sum_{k=1}^K \hat{G}_{\text{IWCV}}^k,$$

where

$$\begin{aligned} \hat{G}_{\text{IWCV}}^k &= \frac{K}{MN} \sum_{d_{m,n}^{\tilde{\pi}} \in \mathcal{D}_{\tilde{\pi}_k}^k} \sum_{n=1}^N \hat{g}(s_{m,n}^{\tilde{\pi}}, a_{m,n}^{\tilde{\pi}}, r_{m,n}^{\tilde{\pi}}; \hat{\theta}_{\text{APDIW}}^k, \mathcal{D}_{\tilde{\pi}_k}^k) \hat{w}_{m,n}. \end{aligned}$$

We estimate the approximation error by the above K -fold IWCV method for all candidate models (in the current setting, a candidate model corresponds to a different value of the flattening parameter ν) and choose the one that minimizes the estimated error:

$$\hat{\nu}_{\text{IWCV}} = \arg \min_{\nu} \hat{G}_{\text{IWCV}}.$$

One may think that, for model selection, \hat{G}_{APDIW} could be directly used, instead of \hat{G}_{IWCV} . However, it can be proven that \hat{G}_{APDIW} is heavily biased (or in other words, *over-fitted*) since the same training samples are used twice for learning parameters and estimating the approximation error [16]. On the other hand, we can prove that \hat{G}_{IWCV} is an almost unbiased estimator of G , where ‘almost’ comes from the fact that the number of training samples is reduced due to data splitting in the cross-validation procedure [16]. Note that ordinary CV (without importance weight) is heavily biased due to the off-policy setup.

B. Numerical Examples

In order to illustrate how IWCV works, we use the same numerical examples as Section IV-B. Fig.3 depicts the approximation error estimated by 5-fold IWCV averaged over 50 trials as a function of the flattening parameter ν . The graphs show that IWCV nicely captures the trend of the true approximation error for all three cases (cf. Fig.2).

Fig.4 describes, as a function of the number M of episodes, the average true approximation error obtained by NIW (APDIW with $\nu = 0$), PDIW (APDIW with $\nu = 1$), APDIW+IWCV ($\nu \in \{0.0, 0.1, \dots, 0.9, 1.0\}$ is selected in each trial using 5-fold IWCV). This result shows that the improvement of the performance by NIW saturates when $M \geq 30$, implying that the bias caused by NIW is not negligible. The performance of PDIW is worse than NIW when $M \leq 20$, which is caused by the large variance of PDIW. On the other hand, APDIW+IWCV consistently gives

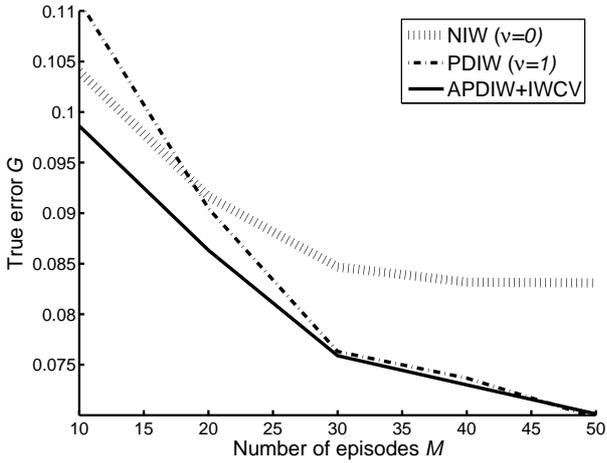


Fig. 4. Average true approximation error G over 50 trials obtained by NIW ($\nu = 0$), PDIW ($\nu = 1$), APDIW+IWCV (ν is chosen by IWCV) in the 10-state chain-walk MDP.

good performance for all M , illustrating the strong adaptation ability of the proposed method.

VI. SAMPLE-REUSE POLICY ITERATION (SRPI)

So far, we have only used our proposed APDIW+IWCV method in the context of policy evaluation. In this section, we extend the method to the full policy-iteration setup.

A. Algorithm

Let us denote the policy at the l -th iteration by π_l and the maximum number of iterations by L . In general policy-iteration methods, new data samples \mathcal{D}^{π_l} are collected following the new policy π_l during the policy evaluation step. Thus, previously-collected data samples $\{\mathcal{D}^{\pi_1}, \mathcal{D}^{\pi_2}, \dots, \mathcal{D}^{\pi_{l-1}}\}$ are not used:

$$\pi_1 \xrightarrow{E:\{\mathcal{D}^{\pi_1}\}} \widehat{Q}^{\pi_1} \xrightarrow{I} \pi_2 \xrightarrow{E:\{\mathcal{D}^{\pi_2}\}} \widehat{Q}^{\pi_2} \xrightarrow{I} \dots \xrightarrow{I} \pi_L,$$

where $E : \{\mathcal{D}\}$ indicates policy evaluation using the data sample \mathcal{D} . It would be more cost-efficient if we could reuse all previously-collected data samples to perform policy evaluation with a growing dataset as:

$$\pi_1 \xrightarrow{E:\{\mathcal{D}^{\pi_1}\}} \widehat{Q}^{\pi_1} \xrightarrow{I} \pi_2 \xrightarrow{E:\{\mathcal{D}^{\pi_1}, \mathcal{D}^{\pi_2}\}} \widehat{Q}^{\pi_2} \xrightarrow{I} \dots \xrightarrow{I} \pi_L.$$

Reusing previously collected data samples turns this into an *off-policy* scenario as the previous policies and the current policy are different unless the current policy has converged to the optimal one. Here, we propose using APDIW+IWCV in policy iteration. For this purpose, we extend the definition of $\widehat{G}_{\text{APDIW}}^l$ so that multiple sampling policies $\{\pi_1, \pi_2, \dots, \pi_l\}$ are taken into account:

$$\widehat{\theta}_{\text{APDIW}}^l \equiv \arg \min_{\theta} \widehat{G}_{\text{APDIW}}^l,$$

$$\widehat{G}_{\text{APDIW}}^l \equiv \frac{1}{LMN} \sum_{l'=1}^l \sum_{m=1}^M \sum_{n=1}^N \widehat{g}(s_{m,n}^{\pi_{l'}}, a_{m,n}^{\pi_{l'}}; \theta, \{\mathcal{D}^{\pi_{l'}}\}_{l'=1}^l) \left(\frac{\prod_{n'=1}^n \pi_l(a_{m,n'}^{\pi_{l'}} | s_{m,n'})}{\prod_{n'=1}^n \pi_{l'}(a_{m,n'}^{\pi_{l'}} | s_{m,n'})} \right)^{\nu_l}, \quad (9)$$

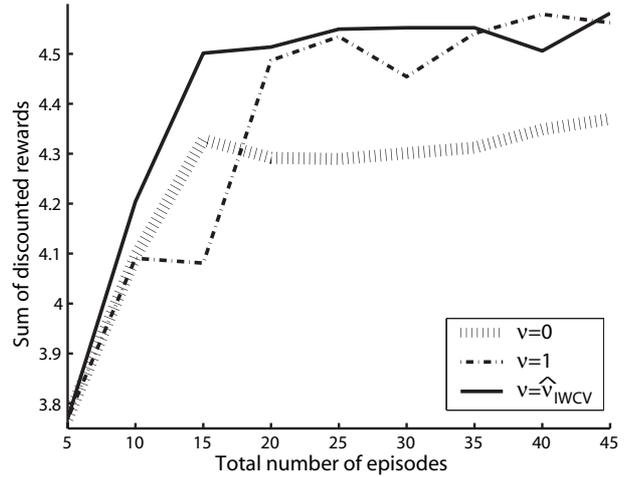


Fig. 5. The performance of policies learned in three scenarios: $\nu = 0$, $\nu = 1$, and SRPI (ν is chosen by IWCV) in the 10-state chain-walk problem.

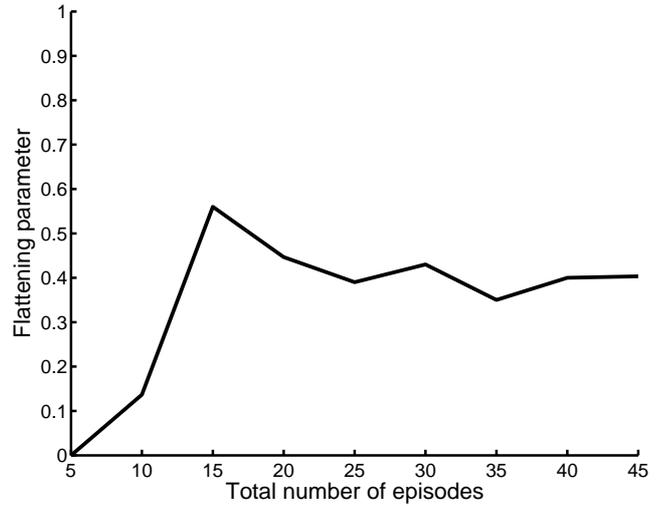


Fig. 6. Average flattening parameter used by SRPI over 30 trials as a function of total number of episodes in the 10-state chain-walk problem.

where $\widehat{G}_{\text{APDIW}}^l$ is the approximation error estimated at the l -th policy evaluation using APDIW. The flattening parameter ν_l is chosen based on IWCV before performing policy evaluation. We call this *sample-reuse policy iteration* (SRPI).

B. Numerical Examples

We will use the system from the numerical examples in Section V-B also in the policy iteration context in order to illustrate how SRPI works. We consider three scenarios: ν is fixed at 0, ν is fixed at 1, and SRPI where ν is chosen by IWCV. The agent collects samples \mathcal{D}^{π_l} at every policy iteration following the current policy π_l , and computes $\widehat{\theta}_{\text{APDIW}}^l$ from all collected samples $\{\mathcal{D}^{\pi_1}, \mathcal{D}^{\pi_2}, \dots, \mathcal{D}^{\pi_l}\}$ using Eq.(9). We use Gaussian kernels defined by Eq.(8) except that kernel centers $\{c_1, c_2, c_3\}$ are randomly selected from the state space \mathcal{S} every trial. The initial policy π_1 is chosen randomly and policy improvement is carried out by the Gibbs policy (2) with $\tau = 2l$.

Fig.5 depicts the average sum of discounted rewards over 30 trials when $M = 5$ with a fixed number of steps ($N =$

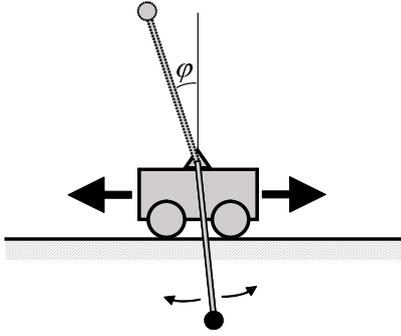


Fig. 7. Illustration of inverted pendulum.

Parameter	Value
Mass of the cart, W	$8[kg]$
Mass of the rod, w	$2[kg]$
Length of the rod, d	$0.5[m]$
Simulation time step, Δt	$0.1[s]$

Fig. 8. Parameters used in the inverted-pendulum simulation.

10). The graphs show that SRPI provides stable and fast learning of policies while the performance improvement of policies learned with $\nu = 0$ saturates in early iterations. The method with $\nu = 1$ can improve policies well but its progress tends to be behind SRPI.

Fig.6 depicts the average value of the flattening parameter used in SRPI as a function of the total number of episodes. The graphs show that the value of the flattening parameter chosen by IWCV tends to rise in the beginning and go down later. At first sight, this does not agree with the general trend of preferring a low-variance estimator in early stages and preferring a low-bias estimator later—but his result is still consistent with the general trend. When the sum of discounted rewards increase rapidly (the total numbers of episodes is up to 15 in Fig.6), the value of the flattening parameter increases (see Fig.5). After that, the sum of discounted rewards does not increase anymore (see Fig.5) since the policy iteration has already been converged. Then it is natural to prefer a smaller flattening parameter (Fig.6) since the sample selection bias becomes mild after convergence.

These results show that SRPI can effectively reuse previously-collected samples by appropriately tuning the flattening parameter according to the condition of data samples, policies etc.

VII. EXPERIMENTS

In this section, we evaluate the performance of our proposed method SRPI in swing-up inverted pendulum problem. We consider the task of *swing-up inverted pendulum* [3] illustrated in Fig.7, consisting of a rod hinged at the top of a cart. The goal of the task is to swing the rod up by moving the cart. We have three actions: applying positive force $+50 [kg \cdot m/s^2]$ to the cart to move right, negative force -50 to move left, and zero force to just coast. That is, the action

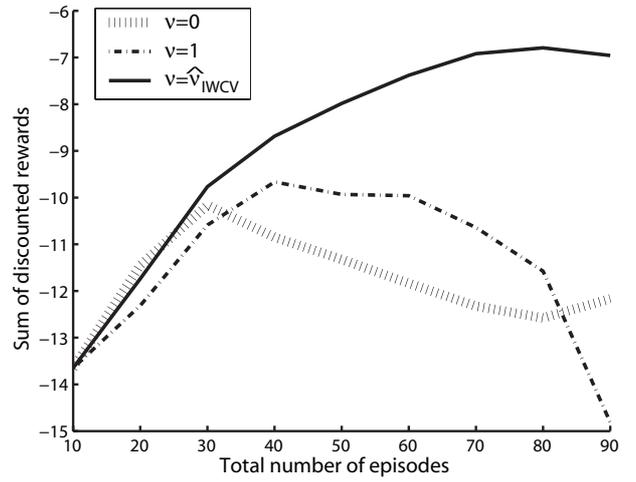


Fig. 9. The results of sample-reuse policy iteration in the task of swing-up inverted-pendulum.

space \mathcal{A} is discrete and described by

$$\mathcal{A} = \{50, -50, 0\} [kg \cdot m/s^2].$$

We note that the force itself is not strong enough to swing the rod up; so the cart needs to be moved back and forth several times to swing the rod up. The state space \mathcal{S} is continuous and consists of the angle $\varphi [rad]$ ($\in [0, 2\pi]$) and the angular velocity $\dot{\varphi} [rad/s]$ ($\in [-\pi, \pi]$)—thus, a state s is described by two-dimensional vector $s = (\varphi, \dot{\varphi})^\top$. Fig.8 shows the parameter setting used in the simulation. The angle φ and angular velocity $\dot{\varphi}$ are updated [18] as follows:

$$\begin{aligned} \varphi_{t+1} &= \varphi_t + \dot{\varphi}_{t+1} \Delta t, \\ \dot{\varphi}_{t+1} &= \dot{\varphi}_t + \frac{9.8 \sin(\varphi_t) - \frac{\alpha w d (\dot{\varphi}_t)^2 \sin(2\varphi_t)}{2} + \alpha \cos(\varphi_t) a_t}{\frac{4l}{3} - \alpha w d \cos^2(\varphi_t)} \Delta t, \end{aligned}$$

where $\alpha = 1/(W + w)$ and a_t is the action ($\in \mathcal{A}$) chosen at time t . We define the reward function $R(s, a, s')$ as

$$R(s, a, s') = \cos(\varphi_{s'}),$$

where $\varphi_{s'}$ denotes the angle φ of state s' .

We use 48 Gaussian kernels with standard deviation $\sigma = \pi$ as basis functions, and arrange kernel centers over the following grid points:

$$\{0, 2/3\pi, 4/3\pi, 2\pi\} \times \{-3\pi, -\pi, \pi, 3\pi\}.$$

That is, the basis functions $\phi(s, a) = \{\phi_1(s, a), \dots, \phi_{16}(s, a)\}$ are set as

$$\phi_{16(i-1)+j}(s, a) = I(a = a^{(i)}) \exp\left(-\frac{\|s - c_j\|^2}{2\sigma^2}\right),$$

for $i = 1, 2, 3$ and $j = 1, 2, \dots, 16$, where

$$c_1 = (0, -3\pi)^\top, c_2 = (0, -\pi)^\top, \dots, c_{12} = (2\pi, 3\pi)^\top.$$

The initial policy $\pi_1(a|s)$ is chosen randomly, and the initial-state probability density $P_1(s)$ is set to be uniform. The agent collects data samples \mathcal{D}^{π_l} ($M = 10$ and $N = 100$) at each policy iteration following the current policy π_l . The

discounted factor is set at $\gamma = 0.95$ and the policy is improved by the Gibbs policy (2) with $\tau = l$.

Fig.9 describes the performance of learned policies. The graph shows that SRPI nicely improves the performance throughout entire policy iteration. On the other hand, the performance when the flattening parameter is fixed at $\nu = 0$ or $\nu = 1$ is not properly improved after the middle of iterations. This result indicates that the flattening parameter is well-adjusted to reuse the previously-collected samples effectively for policy evaluation, and thus SRPI can outperform the other methods.

VIII. CONCLUSIONS AND OUTLOOK

Instability has been one of the critical limitations of importance-sampling techniques, which often makes off-policy methods impractical. To overcome this weakness, we introduced an *adaptive* importance-sampling technique for controlling the trade-off between consistency and stability in value function approximation. We further provided an automatic model selection method for actively choosing the trade-off parameter. We also proposed using the adaptive importance-sampling technique in policy iteration for efficiently reusing previously-collected data samples. The experimental results showed that the proposed method compares favorably with existing approaches.

The method presented in this paper may be extended to policy gradient methods where the need for the sample reuse is even more urgent as small gradient-descent steps result into an under utilization of the data. While importance-sampling has been applied in the settings of policy gradient methods [14], [7], policy gradient methods tend to be unstable when used with standard importance sampling methods [5], [9]—the proposed methods would offer an interesting alternative. Similarly, it can be applied in EM-based methods such as the reward-weighted regression [8], and dynamical-programming based methods such as the policy search by dynamic programming [1].

REFERENCES

- [1] J. A. Bagnell, S. Kakade, A. Y. Ng, and J. Schneider. Policy search by dynamic programming. In *Neural Information Processing Systems 16*, 2003.
- [2] P. D. Bertsekas and J. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, NH, USA, 1996.
- [3] M. Bugeja. Non-linear swing-up and stabilizing control of an inverted pendulum system. In *Proceedings of IEEE Region 8 EUROCON*, pages 437–441, 2003.
- [4] G. S. Fishman. *Monte Carlo: Concepts, Algorithms, and Applications*. Springer-Verlag, Berlin, 1996.
- [5] S. Kakade. A natural policy gradient. In *Neural Information Processing Systems 14*, pages 1531–1538, 2002.
- [6] M. G. Lagoudakis and R. Parr. Least-squares policy iteration. *Journal of Machine Learning Research*, 4(Dec):1107–1149, 2003.
- [7] C. R. Peshkin, L. Shelton. Learning from scarce experience. In *Proceedings of International Conference on Machine Learning*, pages 498–505, 2002.
- [8] J. Peters and S. Schaal. Reinforcement learning by reward-weighted regression for operational space control. In *Proceedings of the International Conference on Machine Learning*, 2007.
- [9] J. Peters and S. Schaal. Reinforcement learning of motor skills with policy gradients. *Neural Networks*, 21:682–97, 2008.

- [10] D. Precup, R. S. Sutton, and S. Dasgupta. Off-policy temporal-difference learning with function approximation. In *Proceedings of International Conference on Machine Learning*, pages 417–424, 2001.
- [11] D. Precup, R. S. Sutton, and S. Singh. Eligibility traces for off-policy policy evaluation. In *Proceedings of International Conference on Machine Learning*, pages 759–766, 2000.
- [12] C. R. Rao. *Linear Statistical Inference and Its Applications*. Wiley, New York, 1973.
- [13] R. Schoknecht. Optimality of reinforcement learning algorithms with linear function approximation. In *Neural Information Processing Systems 15*, pages 1555–1562, 2003.
- [14] C. R. Shelton. Policy improvement for pomdps using normalized importance sampling. In *Proceedings of Uncertainty in Artificial Intelligence*, pages 496–503, 2001.
- [15] H. Shimodaira. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of Statistical Planning and Inference*, 90(2):227–244, 2000.
- [16] M. Sugiyama, M. Krauledat, and K.-R. Müller. Covariate shift adaptation by importance weighted cross validation. *Journal of Machine Learning Research*, 8(May):985–1005, 2007.
- [17] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, MA, USA, 1998.
- [18] H. O. Wang, K. Tanaka, and M. F. Griffin. An approach to fuzzy control of nonlinear systems: Stability and design issues. *IEEE Transactions on Fuzzy Systems*, pages 14–23, 1996.